

SOLUTION TO THE HENCKELL–RHODES PROBLEM: FINITE F -INVERSE COVERS DO EXIST

K. AUINGER, J. BITTERLICH AND M. OTTO

ABSTRACT. For a finite connected graph \mathcal{E} with edge set E , a finite E -generated group G is constructed such that the set of relations $p = 1$ satisfied by G (with p a word over $E \cup E^{-1}$) is closed under deletion of generators (i.e. edges). As a consequence, every element $g \in G$ admits a unique minimal set $C(g)$ of edges (the *content* of g) needed to represent g as a word over $C(g) \cup C(g)^{-1}$. The crucial property of the group G is that connectivity in the graph \mathcal{E} is reflected in G in the following sense: if a word p forms a path $u \rightarrow v$ in \mathcal{E} then there exists a G -equivalent word q which also forms a path $u \rightarrow v$ and uses only edges from their common content; in particular, the content of the corresponding group element $[p]_G = [q]_G$ spans a connected subgraph of \mathcal{E} containing the vertices u and v . Therefore, the construction provides another instance of how certain features of free groups can be “approximated” or “simulated” in finite groups. As an application it is shown that every finite inverse monoid admits a finite F -inverse cover. This solves a long-standing problem of Henckell and Rhodes.

1. INTRODUCTION

In the influential paper [11], Henckell and Rhodes stated a series of conjectures and two problems. The paper was concerned with the celebrated question whether every finite block group M (a monoid in which every von Neumann regular element admits a unique inverse) is a quotient of a submonoid of the power monoid $\mathfrak{P}(G)$ of some finite group G . Henckell and Rhodes presented an affirmative answer to the question modulo some conjecture, namely about the structure of pointlike sets; a subset X of a finite monoid M is *pointlike* (with respect to groups) if and only if in every subdirect product $T \subseteq M \times G$ of M with a finite group G there exists an element $g \in G$ with $X \times \{g\} \subseteq T$ (that is, all elements of X *relate* to some point $g \in G$.) All stated conjectures concerned various aspects of pointlike sets. For example, for inverse monoids M the conjecture was that a subset of M is pointlike if and only if it admits an upper bound with respect to the natural partial order of M . Shortly after, all stated conjectures and one of the two problems (about *liftable tuples* of monoids) were verified respectively solved by Ash in his celebrated paper [3]. The importance of the latter paper went

2020 *Mathematics Subject Classification*. 20M18, 20F05, 20F10, 20F65, 20B99, 05C25.

Key words and phrases. F -inverse monoid, F -inverse cover, permutation group, action graph, Cayley graph.

beyond its immediate task as in the following years interesting and deep connections with the profinite topology of the free group [21] and model theory [12] have been revealed and studied [1, 2]. Yet the second problem stated, which was called by the authors a “stronger form of the pointlike conjecture for inverse monoids”, was not solved in Ash’s paper and has since then attracted considerable attention [14, 25, 26, 5, 24, 23, 8]. It asked:

Problem 1.1. Does every finite inverse monoid admit a finite F -inverse cover?

An inverse monoid is F -inverse if every congruence class of the smallest group congruence σ of F admits a greatest element (with respect to the natural partial order) and an inverse monoid F is a *cover* of an inverse monoid M if there exists a surjective, idempotent separating homomorphism $F \rightarrow M$.

The second author was the first to understand that Problem 1.1 admits a positive solution. In his paper [9] and his dissertation [10] he presented a proof which strongly relied on a result of the third author [17, 18] about the existence of certain finite groupoids. Later, some flaws were discovered in [17, 18] which, however, have been fixed in the meantime [19]. The intention of the present paper is to give a complete and self-contained presentation of the solution to Problem 1.1 (up to classical results on inverse monoids), which is based on the ideas and proofs of [19] but is in a sense tailored for what is needed in the present context and presented in a language which (hopefully) makes it more accessible to the semigroup community. The paper is organised as follows: Section 2 collects prerequisites from inverse monoids, graphs and a proof that the existence of certain groups yields a positive solution of Problem 1.1. Section 3 introduces the main graph-theoretic tools while Section 4 presents two crucial technical results. Finally, in Section 5 we present a construction of the required groups. This construction intends to “reflect the geometry” of a given finite graph.

2. INVERSE MONOIDS

2.1. Preliminaries. 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 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. By the Wagner–Preston Theorem [14, Chapter 1, Theorem 1], inverse monoids may as well be characterized as monoids of partial injective mappings 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.

From basic facts of universal algebra it follows that every inverse monoid M admits a smallest congruence such that the corresponding quotient structure is a group. This congruence is usually denoted σ and it can be characterised as the smallest congruence on M that identifies all idempotents of M with each other. Another way to characterise this congruence is this: two elements $x, y \in M$ are σ -related if and only if $xe = ye$ for some idempotent e of M (and this is equivalent to $fx = fy$ for some idempotent f of M).

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 injective functions then the idempotents of M are exactly the restrictions of the identity function and 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 and inversion of M where the latter means that $x \leq y$ implies $x^{-1} \leq y^{-1}$. In terms of the natural order, the congruence σ can be characterised as the smallest congruence for which the natural order on the quotient is the identity relation, and, likewise as the smallest congruence that identifies every pair of \leq -comparable elements. This leads to yet another description of σ : two elements x and y are σ -related if and only if they admit a common lower bound with respect to \leq . For further information on inverse monoids the reader is referred to the monographs by Petrich [20] and Lawson [14].

An inverse monoid is *F-inverse* if every σ -class possesses a greatest element with respect to \leq . For recent developments concerning the systematic study of *F-inverse* monoids and their relevance in various contexts the reader is referred to [7] and the literature cited there. An *F-inverse* monoid F is an *F-inverse cover* of the inverse monoid M if there exists a surjective idempotent separating homomorphism $F \rightarrow M$. As mentioned in the introduction, it has been an outstanding open problem whether every finite inverse monoid M admits a finite *F-inverse* cover. In order to formulate the following very useful result [20, Theorem VII.6.11] we need the concept of *premorphisms*: for inverse monoids M and N , a mapping $\psi: M \rightarrow N$ is a *premorphisms* if $\psi(1) = 1$, $\psi(m^{-1}) = (\psi m)^{-1}$ and $(\psi m)(\psi n) \leq \psi(mn)$ for all $m, n \in M$.

Theorem 2.1. *Let H be a group and M be an inverse monoid; if $\psi: H \rightarrow M$ is a *premorphisms* such that, for every $m \in M$, there exists $h \in H$ with $m \leq \psi(h)$, then the subdirect product*

$$F := \{(h, m) \in H \times M : m \leq \psi(h)\}$$

is an F-inverse cover of M . Conversely, every F-inverse cover of M can be so constructed.

The following is an easy observation.

Observation 2.2. *Suppose that $\psi: H \rightarrow M$ is as in Theorem 2.1 and $\pi: M \rightarrow N$ is a surjective homomorphism with N an inverse monoid; then the composition $\pi \circ \psi: H \rightarrow N$ is a premorphism which also satisfies the condition of Theorem 2.1.*

Hence the task for Problem 1.1 is, given a finite inverse monoid N , to find a finite group H which admits a premorphism $H \rightarrow N$ satisfying the condition of Theorem 2.1. Observation 2.2 eases the situation a bit since we need to do so only for a special kind of inverse monoids (in the rôle of M) which we shall describe below (see § 2.4).

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 $(x_1 \cdots x_n)^{-1} = x_n^{-1} \cdots x_1^{-1}$ (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 let 1 denote the *empty 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 $\psi: M \twoheadrightarrow N$. In this situation, M is an *expansion of N* . The special case of A -generated groups will play a significant rôle in this paper.

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. In particular M/σ , the quotient of an A -generated inverse monoid M modulo the smallest group congruence σ , is an A -generated group.

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 -equivalent*, while M avoids the relation $p = q$ if $[p]_M \neq [q]_M$.

2.3. Graphs. In this paper, we consider a very liberal notion of graph structures, admitting multiple directed edges between pairs of vertices and including directed loops at individual vertices. In the literature, such structures are often called *directed multigraphs* or *quivers*. The following formalisation is convenient for our purposes. A *graph* \mathcal{E} is a *two-sorted structure* $(V \cup K; \alpha, \omega, {}^{-1})$ with V its set of *vertices*, K its set of *edges* (disjoint from V), with *incidence functions* $\alpha: K \rightarrow V$ and $\omega: K \rightarrow V$, selecting, for each edge e 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 $\alpha e = u$ and $\omega e = v$ in “geometric” terms as $e: \underset{u}{\bullet} \longrightarrow \underset{v}{\bullet}$ and its inverse $e^{-1}: \underset{u}{\bullet} \longleftarrow \underset{v}{\bullet}$ as “the same edge but traversed in the opposite direction”. 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. An oriented graph \mathcal{E} with set of positive edges E will be denoted as $\mathcal{E} = (V \cup \tilde{E}; \alpha, \omega, {}^{-1})$.

A *subgraph* of the graph \mathcal{E} 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 subset $S \subseteq V \cup K$ generates a unique minimal subgraph $\langle S \rangle$ of \mathcal{E} containing S , which is the *subgraph of \mathcal{E} spanned by S* . An *automorphism* of a graph $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ is a map $\varphi = \varphi_V \cup \varphi_K: V \cup K \rightarrow V \cup K$ with $\varphi_V: V \rightarrow V$, $\varphi_K: K \rightarrow K$ being bijections satisfying for all $e \in K$:

$$\alpha \varphi_K(e) = \varphi_V(\alpha e), \quad \omega \varphi_K(e) = \varphi_V(\omega e), \quad \varphi_K(e^{-1}) = (\varphi_K(e))^{-1}.$$

We note that the second equality is a consequence of the first and third. In the oriented case we require in addition that $\varphi_{\tilde{E}}(E) = E$ and (therefore also) $\varphi_{\tilde{E}}(E^{-1}) = E^{-1}$. A benefit from our definition of a graph as a two-sorted functional rather than a relational structure is that there is no distinction between weak and induced subgraphs and that concepts like homomorphism, congruence and quotient are easier to handle.

Let A be a finite set; a *labelling* of the graph $\mathcal{E} = (V \cup K; \alpha, \omega, {}^{-1})$ by the *alphabet* A (an *A-labelling*, for short) is a mapping $\ell: K \rightarrow \tilde{A}$ respecting the involution: $\ell(e^{-1}) = \ell(e)^{-1}$ for all $e \in K$. The labelling $\ell: K \rightarrow \tilde{A}$ gives rise to an orientation of \mathcal{E} : 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 $K = \tilde{E}$.

We consider A -labelled graphs as structures $(V \cup K; \alpha, \omega, {}^{-1}, \ell, A)$ in their own right. By a *subgraph of an A-labelled graph* we mean just a subgraph with the induced labelling. Morphisms of A -labelled graphs are naturally defined as follows. Let $\mathcal{K} = (V \cup K; \alpha, \omega, {}^{-1}, \ell, A)$ and $\mathcal{L} = (W \cup L; \alpha, \omega, {}^{-1}, \ell, A)$

be A -labelled graphs. A *morphism* $\varphi: \mathcal{K} \rightarrow \mathcal{L}$ of A -labelled graphs is a mapping $\varphi: V \cup K \rightarrow W \cup L$, mapping vertices to vertices and edges to edges, that is compatible with the operations α and $^{-1}$ (and therefore also ω) as well as with the labelling.

A *congruence* Θ on the A -labelled graph $\mathcal{K} = (V \cup K; \alpha, \omega, ^{-1}, \ell, A)$ is an equivalence relation on $V \cup K$ contained in $(V \times V) \cup (K \times K)$ which is compatible with the operations α and $^{-1}$ (therefore also ω) and respects ℓ :

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

and

$$e \Theta f \implies \ell(e) = \ell(f) \text{ for all } e, f \in K.$$

The definition of the quotient graph \mathcal{K}/Θ for a congruence Θ is obvious, and we have the usual Homomorphism Theorem.

A *non-empty path* π in \mathcal{E} 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* if $\alpha\pi = \omega\pi$. 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 graph is connected if any two vertices can be connected by some path. 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). For a path $e_1 \cdots e_k$ in an A -labelled graph \mathcal{E} , its label is $\ell(e_1 \cdots e_k) := \ell(e_1) \cdots \ell(e_k)$ which is a word in \tilde{A}^* .

2.4. Cayley graphs of A -generated groups and the Margolis–Meakin expansion. Given an A -generated group Q we define the *Cayley graph* \mathcal{Q} of Q by the following data; as an A -labelled graph, this graph \mathcal{Q} depends on the underlying assignment function $i_{\mathcal{Q}}$:

- 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

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

The edge (g, a) should be thought of as $\bullet \xrightarrow{a} \bullet$, its inverse as $\bullet \xleftarrow{a^{-1}} \bullet$, where ga stands for $g[a]_{\mathcal{Q}}$. 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 is an element of the acting group Q , g a vertex of \mathcal{Q} and (g, a) an edge of \mathcal{Q} .

We arrive at the important concept of the *Margolis–Meakin expansion* $M(Q)$ of an A -generated group Q [15]. For a given A -generated group Q , the Margolis–Meakin expansion $M(Q)$ consists of all pairs (\mathcal{K}, g) with $g \in Q$ and \mathcal{K} a finite connected subgraph of the Cayley graph \mathcal{Q} of Q containing the vertices 1 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 $(\{1\}, 1)$ and with respect to the assignment function

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

The value of some 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 and having label p ; the natural partial order on $M(Q)$ is given by

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

The Margolis–Meakin expansion plays an important rôle in the theory of inverse semigroups; for its universal property the reader is referred to [15] or [7]. Most relevant for our purpose is the following, which is a consequence of the results of [15].

Theorem 2.3. *Every finite inverse monoid M arises as a quotient of the Margolis–Meakin expansion $M(Q)$ of some finite A -generated group Q for some finite alphabet A .*

Consequently, in order to find, for a finite inverse monoid M , a finite group H with a premorphism $\psi: H \rightarrow M$ satisfying the condition of Theorem 2.1, according to Observation 2.2 it is sufficient to do so for M being the Margolis–Meakin expansion $M(Q)$ of any finite A -generated group Q .

2.5. F -inverse covers. For a given A -generated group Q as above, we now seek to provide an expansion H of Q , which will allow us to use Theorems 2.1 and 2.3 together with Observation 2.2 towards the construction of F -inverse covers, as in Theorem 2.7 below. First we isolate an important property of groups generated by an alphabet.

Definition 2.4 (*X-generated group with content function*). Let X be any alphabet; an X -generated group R has a *content function* C if for every element $g \in R$ there is a unique \subseteq -minimal subset $C(g)$ of X such that g is represented as a product of elements of $C(g)$ and their inverses.

We need to define one further property, which will be crucial towards the construction of a group H admitting a premorphism $\psi: H \rightarrow M(Q)$, for the Margolis–Meakin expansion to satisfy the condition of Theorem 2.1.

Definition 2.5 (*group reflecting the structure of a Cayley graph*). Let Q be an A -generated group with Cayley graph \mathcal{Q} and let $E = Q \times A$ be the set of positive edges of \mathcal{Q} . An E -generated group G reflects the structure of \mathcal{Q} if the following hold.

- (1) The action of Q on E by left multiplication extends to an action of Q on G by automorphisms on the left (denoted $(g, \xi) \mapsto {}^g\xi$ for $g \in Q$ and $\xi \in G$).
- (2) G has a content function C such that, for any $p \in \tilde{E}^*$ which forms a path $g \rightarrow h$ in \mathcal{Q} the following hold:
 - (a) if $C([p]_G) = \emptyset$, that is if $[p]_G = 1$, then $g = h$,
 - (b) if $C([p]_G) \neq \emptyset$, that is if $[p]_G \neq 1$, then there exists a word $q \in \tilde{E}^*$ which also forms a path $g \rightarrow h$ and such that $[p]_G = [q]_G$ and q uses only edges of the content $C([p]_G)$ of $[p]_G$ (and their inverses). In particular, the content $C([p]_G)$ spans a connected subgraph of \mathcal{Q} containing g and h .

Next let Q be an A -generated group and, for $E = Q \times A$, let G be an E -generated group which reflects the structure of the Cayley graph \mathcal{Q} of Q . Since Q acts on G by automorphisms on the left, we can form the semidirect product $G \rtimes Q$, which consists of the set $G \times Q$ endowed with the binary operation

$$(\gamma, g)(\eta, h) := (\gamma \cdot {}^g\eta, gh),$$

inversion

$$(\gamma, g)^{-1} := (g^{-1}\gamma^{-1}, g^{-1})$$

and identity element $(1_G, 1_Q)$. Consider the following A -generated subgroup H of $G \rtimes Q$:

$$H := \langle ([1, a]_G, [a]_Q) : a \in A \rangle \leq G \rtimes Q. \quad (2.1)$$

Similarly to the value $[p]_{M(Q)}$ mentioned above one has that for a word $p \in \tilde{A}^*$ the value of p in H is

$$[p]_H = ([\pi_1^{\mathcal{Q}}(p)]_G, [p]_Q) \quad (2.2)$$

where, again, $\pi_1^{\mathcal{Q}}(p)$ is the unique path in \mathcal{Q} starting at 1 and being labelled p , interpreted as a word over \tilde{E} . In particular, H is an expansion of Q with canonical morphism $([\pi_1^{\mathcal{Q}}(p)]_G, [p]_Q) \mapsto [p]_Q$.

Theorem 2.6. *Let Q be an A -generated group; for $E = Q \times A$, let G be an E -generated group (with content function C) which reflects the structure of the Cayley graph \mathcal{Q} of Q (Definition 2.5), and let H be the group defined by (2.1). Then the mapping*

$$\psi: H \rightarrow M(Q), \quad (\gamma, g) \mapsto \begin{cases} (\{1_Q\}, 1_Q) & \text{if } (\gamma, g) = (1_G, 1_Q) \\ (\langle C(\gamma) \rangle, g) & \text{if } (\gamma, g) \neq (1_G, 1_Q) \end{cases}$$

is a premorphism which satisfies the condition formulated in Theorem 2.1.

Proof. Recall that for $\gamma \in G$, $C(\gamma) = \emptyset$ if and only if $\gamma = 1_G$. The definition of ψ makes sense only if $(1_G, g) \in H$ implies $g = 1_Q$. Let $p \in \tilde{A}^*$ be such that $[p]_H = (1_G, g)$; then $1_G = [\pi_1^{\mathcal{Q}}(p)]_G$ and $\pi_1^{\mathcal{Q}}(p)$ is the path in \mathcal{Q} starting at 1_Q and being labelled p . By (2.2) the terminal vertex of this path is $[p]_Q = g$; but from Definition 2.5 (2a) it follows that this path is closed, hence $[p]_Q = 1_Q$. So, if $[p]_H \neq 1_H$ then $[\pi_1^{\mathcal{Q}}(p)]_G \neq 1_G$ and, by Definition 2.5 (2b), the content $C([\pi_1^{\mathcal{Q}}(p)]_G)$ spans a connected subgraph of \mathcal{Q} containing 1_Q and $[p]_Q$ so that $\psi([p]_H) \in M(Q)$, as required. Definition 2.5, the equalities $C(\gamma^{-1}) = C(\gamma)$, $C({}^g\gamma) = {}^gC(\gamma)$ and the inclusion $C(\gamma \cdot \eta) \subseteq C(\gamma) \cup C(\eta)$ for any $\gamma, \eta \in G$ imply that ψ is a premorphism.

Finally, let $(\mathcal{K}, g) \in M(Q)$ and $p \in \tilde{A}^*$ be such that $[p]_{M(Q)} = (\mathcal{K}, g)$. Then $\mathcal{K} = \langle \pi_1^{\mathcal{Q}}(p) \rangle$, $g = [p]_Q$ and $[p]_H = (\gamma, g)$ where $\gamma = [\pi_1^{\mathcal{Q}}(p)]_G$ (the G -value of the path $\pi_1^{\mathcal{Q}}(p)$). If $C(\gamma) = \emptyset$, that is, $\gamma = 1_G$ then $g = 1_Q$, hence $\psi(\gamma, g) = (\{1_Q\}, 1_Q) \geq (\mathcal{K}, 1_Q) = (\mathcal{K}, g)$. Otherwise, if $C(\gamma) \neq \emptyset$ then every edge of $C([\pi_1^{\mathcal{Q}}(p)]_G)$ belongs to $\langle \pi_1^{\mathcal{Q}}(p) \rangle$, hence

$$\langle C(\gamma) \rangle = \langle C([\pi_1^{\mathcal{Q}}(p)]_G) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(p) \rangle,$$

so that $\psi(\gamma, g) = (\langle C(\gamma) \rangle, g) \geq (\mathcal{K}, g)$, as required. \square

As a consequence of Theorem 2.1, an F -inverse cover of $M(Q)$ can be obtained as a subdirect product of H with $M(Q)$. Observation 2.2 in combination with Theorem 2.3 now implies the result promised in the title of the paper.

Theorem 2.7. *Every finite inverse monoid admits a finite F -inverse cover.*

It can even be shown that the so constructed F -inverse cover of $M(Q)$ is A -generated as an inverse monoid.

Corollary 2.8. *The F -inverse cover F of $M(Q)$ given by Theorem 2.1 as guaranteed by Theorem 2.6 coincides with the subdirect product*

$$S := \{([p]_H, [p]_{M(Q)}) \in H \times M(Q) : p \in \tilde{A}^*\}.$$

In particular, $F = S$ is A -generated as an inverse monoid.

Proof. It is clear that for every $p \in \tilde{A}^*$ the inequality $[p]_{M(Q)} \leq \psi([p]_H)$ holds, hence $S \subseteq F$. Therefore, according to Theorem 2.1, we need to show that every pair $(h, m) \in F$, that is, $(h, m) \in H \times M(Q)$ such that

$m \leq \psi(h)$ belongs to S . Let $(h, m) \neq (1_H, 1_{M(Q)})$ be such a pair and $h = [s]_H$ and $m = [r]_{M(Q)}$ for suitable words $s, r \in \tilde{A}^*$. According to (2.2), $[s]_H = ([\pi_1^{\mathcal{Q}}(s)]_G, [s]_Q)$ where $\pi_1^{\mathcal{Q}}(s)$ is the path $1 \rightarrow [s]_Q$ in \mathcal{Q} starting at 1 and being labelled s . By Definition 2.5 (2b) there exists a path $\pi: 1 \rightarrow [s]_Q$ in \mathcal{Q} such that $[\pi_1^{\mathcal{Q}}(s)]_G = [\pi]_G$ and π uses only edges from the G -content $C([\pi_1^{\mathcal{Q}}(s)]_G)$ (and their inverses). Let $t := \ell(\pi)$ be the label of that path; then $\pi = \pi_1^{\mathcal{Q}}(t)$ and since $\pi: 1 \rightarrow [s]_Q$ we also have $[t]_Q = [s]_Q$. Altogether, $[t]_H = [s]_H = h$; from

$$(\langle C([\pi_1^{\mathcal{Q}}(t)]_G) \rangle, [t]_Q) = \psi([t]_H) = \psi(h) \geq m = (\langle \pi_1^{\mathcal{Q}}(r) \rangle, [r]_Q)$$

we have that $\langle \pi_1^{\mathcal{Q}}(t) \rangle = \langle C([\pi_1^{\mathcal{Q}}(t)]_G) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(r) \rangle$ and $[t]_Q = [r]_Q$. For $q := tr^{-1}r$ this implies that $[q]_H = h$ and $[q]_{M(G)} = m$, hence $(h, m) \in S$. \square

In order to prove Theorem 2.6 and hence Theorem 2.7 it is sufficient to construct, for any finite A -generated group Q and $E = Q \times A$ a finite E -generated group G which reflects the structure of the Cayley graph \mathcal{Q} of Q according to Definition 2.5. The existence of such a group G is guaranteed by the following more general lemma, which is the main result of the paper. For item (1) recall that every automorphism of an oriented graph induces a permutation of its set of positive edges.

Lemma 2.9 (main lemma). *For every finite connected oriented graph $\mathcal{E} = (V \cup \tilde{E}; \alpha, \omega, {}^{-1})$ there exists a finite E -generated group G which has the following properties:*

- (1) *Every permutation of E induced by an automorphism of \mathcal{E} extends to an automorphism of G .*
- (2) *The set of relations $p = 1$ satisfied by G (with $p \in \tilde{E}^*$) is closed under the deletion of generators and thus G has a content function C (Proposition 3.5).*
- (3) *For any word $p \in \tilde{E}^*$ which forms a path $u \rightarrow v$ in \mathcal{E} (with u and v not necessarily distinct vertices of \mathcal{E}) the following hold:*
 - (a) *if $C([p]_G) = \emptyset$ then $u = v$,*
 - (b) *if $C([p]_G) \neq \emptyset$ then there exists a word $q \in \tilde{E}^*$ with $[p]_G = [q]_G$ such that q also forms a path $u \rightarrow v$ and q only uses edges from the content $C([p]_G)$ (and their inverses). In particular, $C([p]_G)$ spans a connected subgraph of \mathcal{E} containing u and v .*

Remark 2.10. The free group generated by E obviously enjoys properties (1)–(3) of Lemma 2.9. Hence, the main result of the paper is another instance of when the behaviour of a free group can be “simulated” or “approximated” by a finite group [1, 2, 3, 12, 16].

The remainder of the paper is devoted to proving Lemma 2.9. This requires quite a bit of work. It will be accomplished in Section 5. In order to achieve this goal we introduce several graph-theoretic constructions which

will be presented in Sections 3 and 4. The results in these two sections are of a more general nature and are of independent interest.

3. TOOLS

In this section we introduce some graph-theoretic constructions, which later will enable the construction of a group G as mentioned above. The group itself will be realised as a permutation group defined by its *action graph*. It is a well-established approach to construct finite A -generated groups which avoid certain unwanted relations, to proceed as described in the following. First encode the relations in a finite A -labelled directed graph \mathcal{X} — the set of unwanted relations will usually be infinite, but must in some sense be regular (recognisable by a finite automaton). Next take a quotient \mathcal{X}/\equiv of \mathcal{X} which guarantees that the edge labels from A induce partial injective mappings on the vertex set. Finally form some *completion* $\overline{\mathcal{X}/\equiv}$ of \mathcal{X}/\equiv , through extending the partial injective mappings to total permutations of the vertex set of \mathcal{X}/\equiv or of some finite superset. The letters $a \in A$ then act as permutations on the finite set of vertices of $\overline{\mathcal{X}/\equiv}$ and one gets a finite permutation group that avoids the unwanted relations. A meanwhile classical application of this approach is Stallings' proof of Hall's Theorem that every finitely generated subgroup of a free group F is closed in the profinite topology of F [22]. Many more examples can be found in [13, 6] and elsewhere. In his paper [3] Ash definitely developed some mastership of arguments of this kind. Independently, the third author has suggested a considerable refinement of this approach [16]. He proposed a construction which is inductive on the subsets of the generating set A in the sense that the k th group G_k satisfies/avoids all relations $p = 1$ in at most k letters that should be satisfied/avoided by the final group G . In the step $G_k \rightsquigarrow G_{k+1}$ not only new relations $p = 1$ in more than k letters are added which are to be avoided (by adding components to the graph which defines G_k) but, at the same time, the relations in at most k letters must be preserved. The motivation for this approach has come from some relevant applications to hypergraph coverings and finite model theory [16]. The constructions in this section and the results of the next section are of this flavour and are taken from the third author's [19].

3.1. E -graphs and E -groups. We slightly change perspective: since the edges of the graph \mathcal{E} of Lemma 2.9 are the letters of the labelling alphabet we now denote the labelling alphabet by E . An E -labelled graph is an E -graph if every vertex u has, for every label $a \in \tilde{E}$, 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* [13] or *inverse automaton* [4, 6], to mention just two. In an E -graph \mathcal{K} , for every word $p \in \tilde{E}^*$ and every vertex u there is at most one path $\pi = \pi_u^{\mathcal{K}}(p)$ with initial vertex $\alpha\pi = u$ and label $\ell(\pi) = p$. For a path π in \mathcal{K} with initial vertex u , terminal vertex v and label

$p \in \tilde{A}^*$ (for $A \subseteq E$) we write $u \xrightarrow{p} v$ or $p: u \rightarrow v$ and we call π an A -path $u \rightarrow v$; the vertices u and v are A -connected in \mathcal{K} . The A -component of a vertex v of the E -graph \mathcal{K} , denoted $v\mathcal{K}[A]$, is the subgraph of \mathcal{K} spanned by all paths in \mathcal{K} having initial vertex v and whose labels are in \tilde{A}^* . A labelled graph \mathcal{K} is called *complete* or a *group action graph* (also called *permutation automaton*) if every vertex u has, for every label $a \in \tilde{E}$ exactly one edge f with initial vertex $\alpha f = u$ and label $\ell(f) = a$; in this case, for every word $p \in \tilde{E}^*$ and every vertex u there exists exactly one path $\pi = \pi_u^{\mathcal{K}}(p)$ starting at u and having label 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 ; then, for every $p \in \tilde{E}^*$, the mapping $[p]: V \rightarrow V$, $u \mapsto u \cdot p$ is a permutation of the vertex set V of \mathcal{K} . Thus the involutory monoid \tilde{E}^* acts on V by permutations on the right. The permutation group

$$\mathcal{T}(\mathcal{K}) := \{[p]: p \in \tilde{E}^*\} \quad (3.1)$$

obtained this way, is called the *transition group* $\mathcal{T}(\mathcal{K})$ of the graph \mathcal{K} . This transition group $\mathcal{T}(\mathcal{K})$ is an E -generated group (E -group for short) in a natural way, the letter $e \in \tilde{E}$ induces the permutation $[e]$ which maps every vertex u to the terminal vertex $\omega\pi_u(e)$ of the edge $\pi_u(e)$ which is the unique edge with initial vertex u and label e . Note that this edge may be a loop edge for every vertex u (so $[e]$ might be the identity element of $\mathcal{T}(\mathcal{K})$). Moreover, it may happen that distinct letters $e \neq f \in \tilde{E}$ induce the same permutation.

A crucial fact concerning the transition group $G = \mathcal{T}(\mathcal{K})$ is the following: for every connected component \mathcal{C} of \mathcal{K} and every vertex u of \mathcal{C} there is a unique surjective graph morphism $\varphi_u: \mathcal{G} \rightarrow \mathcal{C}$ from the Cayley graph \mathcal{G} of G onto \mathcal{C} for which $\varphi_u(1) = u$; we call φ_u the *canonical morphism* $\mathcal{G} \rightarrow \mathcal{C}$ with respect to u ; occasionally we shall leave the vertex u undetermined and shall speak of *some* canonical morphism $\mathcal{G} \rightarrow \mathcal{C}$. For easy reference we give a name to this phenomenon.

Definition 3.1. The Cayley graph \mathcal{G} of an E -group G covers a complete, connected E -graph \mathcal{C} if there is a canonical morphism $\varphi: \mathcal{G} \rightarrow \mathcal{C}$.

An E -graph $(V \cup \tilde{E}; \alpha, \omega, {}^{-1})$ is *weakly complete* if, for every letter $a \in \tilde{E}$, the partial injective mapping on V induced by a is a permutation on its domain; in other words, provided that the graph is finite, the subgraph spanned by all edges with label a is a disjoint union of cycle graphs (a -cycles). For every weakly complete graph \mathcal{K} we denote by $\overline{\mathcal{K}}$ its *trivial completion*, that is, the complete graph obtained by adding, for every $a \in \tilde{E}$, a loop edge with label a to every vertex not already contained in an a -cycle of \mathcal{K} .

3.2. k -retractable groups, content function and k -stable expansions.

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 . Let G be an E -group; for every $A \subseteq E$ let $G[A]$ be the A -generated subgroup of G .

Definition 3.2. An E -group G is *retractable* if, for all words $p, q \in \widetilde{E}^*$ and every letter $a \in E$ the following holds: ¹

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

Moreover, G is *A-retractable* if $G[A]$ is retractable (as an A -group), and, for $k \leq |E|$, G is *k-retractable* if G is A -retractable for every $A \subseteq E$ with $|A| = k$.

Of course, k -retractability implies l -retractability for all $l \leq k$, and every group is 1-retractable. Retractability of an E -group G means that for every subset $A \subseteq E$ the mapping

$$E \rightarrow E \cup \{1\}, a \mapsto \begin{cases} a & \text{if } a \in A \\ 1 & \text{if } a \notin A \end{cases}$$

extends to an endomorphism ψ_A of G , which in fact is a retract endomorphism onto $G[A]$ (the image of ψ_A is $G[A]$ and its restriction to $G[A]$ is the identity mapping). For an E -group G and $A \subseteq E$ we denote the Cayley graph of $G[A]$, considered as an A -graph, by $\mathcal{G}[A]$; this graph is weakly complete as an E -graph and, as above, we denote its trivial completion by $\overline{\mathcal{G}[A]}$. In light of the connection with retract endomorphisms we see the following.

Proposition 3.3. *An E -group G is retractable if and only if its Cayley graph \mathcal{G} covers $\overline{\mathcal{G}[A]}$ for every $A \subseteq E$.*

Proof. Suppose that G is retractable and $A \subseteq E$. The retract endomorphism ψ_A is a canonical morphism $\psi_A: G \rightarrow G[A]$ if $G[A]$ is considered as an E -group with all $e \in E \setminus A$ being identity generators. Its Cayley graph with respect to E coincides with $\overline{\mathcal{G}[A]}$. It follows that there is a canonical graph morphism $\mathcal{G} \rightarrow \overline{\mathcal{G}[A]}$, that is, \mathcal{G} covers $\overline{\mathcal{G}[A]}$.

Suppose conversely that for every $A \subseteq E$ there is a canonical graph morphism $\mathcal{G} \rightarrow \overline{\mathcal{G}[A]}$. We note that this morphism must be injective when restricted to $\mathcal{G}[A]$ (considered as a subgraph of \mathcal{G}). Let $p \in \widetilde{E}^*$, $a \in E$ and suppose that $[p]_G = 1$. Then p labels a closed path $\pi_1^{\mathcal{G}}(p)$ at 1 in \mathcal{G} . Let $B = E \setminus \{a\}$. The canonical morphism $\mathcal{G} \rightarrow \overline{\mathcal{G}[B]}$ maps the path $\pi_1^{\mathcal{G}}(p)$ to the path $\pi_1^{\overline{\mathcal{G}[B]}}(p)$ which is also closed. The paths $\pi_1^{\overline{\mathcal{G}[B]}}(p)$ and $\pi_1^{\overline{\mathcal{G}[B]}}(p_{a \rightarrow 1})$ traverse the same edges except loop edges labelled $a^{\pm 1}$, and therefore visit the same vertices. So $\pi_1^{\overline{\mathcal{G}[B]}}(p_{a \rightarrow 1})$ is also closed, and as it runs entirely in $\mathcal{G}[B]$, it follows that $[p_{a \rightarrow 1}]_{\mathcal{G}[B]} = 1$ and therefore $[p_{a \rightarrow 1}]_G = 1$. \square

For a word $p \in \widetilde{E}^*$ the *content* $\text{co}(p)$ is the set of all letters $a \in E$ for which a or a^{-1} occurs in p . The importance of retractable E -groups for our purpose comes from the fact that such E -groups admit a *content function* (Definition 2.4). Indeed, assume that G is retractable. Then, for $p, q \in \widetilde{E}^*$ and $a \in E$ the equality $[p]_G = [q]_G$ implies $[p_{a \rightarrow 1}]_G = [q_{a \rightarrow 1}]_G$. Suppose

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

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 delete (without changing its value $[p]_G$) every letter in a word p which does not occur in every other representation q of the group element $[p]_G$. This leads to the following definition.

Definition 3.4. Let G be a retractable E -group and $g \in G$. The *content* $C(g)$ of g is

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

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

Proposition 3.5. *Every retractable group has a content function.*

In case G is retractable, for any two subsets $A, B \subseteq E$ we have

$$G[A] \cap G[B] = G[A \cap B]. \quad (3.2)$$

Groups satisfying this condition for all $A, B \subseteq E$ have been called *2-acyclic* by the third author in [16, 19]: condition (3.2) rules out patterns as on the left-hand side of Figure 1 where g would belong to $G[A] \cap G[B]$ but not to $G[A \cap B]$, and the cosets $G[A]$ and $G[B]$ form a non-trivial 2-cycle. In other words, condition (3.2) implies that the intersection of two cosets $gG[A]$ and $hG[B]$ in G is either empty or is a coset of the form $kG[A \cap B]$. In terms of connectivity in the Cayley graph \mathcal{G} of G this means that, if two vertices u and v are connected by an A -path as well as by a B -path, then there is even an $(A \cap B)$ -path $u \rightarrow v$; this point of view will be frequently used in the paper.

But indeed, retractable groups also avoid patterns as on the right-hand side of Figure 1. In the terminology of [16, 19], they are even *3-acyclic*. This means that, for all $A, B, C \subseteq E$ and all $g, h, k \in G$ the following holds:

$$\begin{aligned} gG[A] = hG[A], \quad hG[B] = kG[B] \quad \text{and} \quad kG[C] = gG[C] \\ \implies hG[A \cap B] \cap kG[B \cap C] \cap gG[C \cap A] \neq \emptyset, \end{aligned} \quad (3.3)$$

as we shall see in passing, in connection with the proof of Lemma 3.11 below.

Remark 3.6. Retractable E -groups are 2- and 3-acyclic in the sense of satisfying conditions 3.2 and 3.3, meaning that their Cayley graphs do not admit connectivity patterns of cosets as in Figure 1.

Definition 3.7. For $A \subseteq E$, an expansion $H \twoheadrightarrow G$ of E -groups is *A-stable* if the canonical morphism is injective when restricted to $H[A]$; it is *k-stable* (for $k < |E|$) if it is A -stable for every k -element subset A of E .

We arrive at our first basic construction. Here and in the following we use \sqcup and \bigsqcup to denote the disjoint union of graphs; recall the definition of the transition group of a complete graph (3.1).




FIGURE 1

Theorem 3.8. *Let \mathcal{X} be a complete E -graph, $1 \leq k < |E|$ and suppose that the transition group $G = \mathcal{T}(\mathcal{X})$ is k -retractable. Then the transition group*

$$H := \mathcal{T} \left(\mathcal{X} \sqcup \bigsqcup_{|C|=k} \overline{\mathcal{G}[C]} \right)$$

is $(k + 1)$ -retractable and is a k -stable expansion of G . Moreover, every k -stable expansion of H is also $(k + 1)$ -retractable.

Proof. We first show that H is a k -stable expansion of G . So, let $p \in \tilde{E}^*$ be a word with $|\text{co}(p)| \leq k$ and suppose that $[p]_G = 1$. We need to show that $[p]_H = 1$. In order to do so it is sufficient to show that, for every vertex v in $\mathcal{X} \sqcup \bigsqcup_{|C|=k} \overline{\mathcal{G}[C]}$ the path $\pi_v(p)$ which starts at v and has label p is a cycle. This is obvious for every $v \in \mathcal{X}$ and $v \in \overline{\mathcal{G}[A]}$ when A is a set of k letters for which $p \in \tilde{A}^*$. So, let $B \subseteq E$ with $|B| = k$ and suppose that $p \notin \tilde{B}^*$, which means that at least one element of the content of p does not belong to B , and let v be a vertex of $\overline{\mathcal{G}[B]}$. Let p' be the word obtained from p by deletion of all letters from $\text{co}(p) \setminus B$. Since G is k -retractable, we have $[p']_G = 1$ and hence also $[p']_{G[B]} = 1$ since p' contains only letters from B . It follows that the path $\pi_v^{\mathcal{G}[B]}(p')$ is closed and hence so is $\pi_v^{\overline{\mathcal{G}[B]}}(p')$. Since the paths $\pi_v^{\overline{\mathcal{G}[B]}}(p)$ and $\pi_v^{\overline{\mathcal{G}[B]}}(p')$ meet exactly the same vertices — the two paths differ only in loop edges labelled by letters from $\text{co}(p) \setminus B$ — it follows that $\pi_v^{\overline{\mathcal{G}[B]}}(p)$ is also closed. Altogether, $[p]_H = 1$ and the expansion $H \twoheadrightarrow G$ is k -stable.

Let $K \twoheadrightarrow H$ be a k -stable expansion; then the expansion $K \twoheadrightarrow G$ is also k -stable. We show that K is $(k + 1)$ -retractable, this then applies also to $K = H$. So let $A \subseteq E$ with $|A| = k + 1$; according to Proposition 3.3 it suffices to show that for every subset $B \subsetneq A$ there is a canonical morphism $\mathcal{K}[A] \twoheadrightarrow \overline{\mathcal{K}[B]}$ where the latter completion is with respect to $\mathcal{K}[B]$ as an A -graph, that is, loop edges labelled by letters from $A \setminus B$ (and their inverses) are added to all vertices of $\mathcal{K}[B]$. From the definition of H and the assumption on K it follows that there is a canonical morphism $\mathcal{K} \twoheadrightarrow \overline{\mathcal{G}[B]}$. But $K \twoheadrightarrow G$ is k -stable, hence $K[B] \cong G[B]$ and therefore also $\overline{\mathcal{K}[B]} \cong \overline{\mathcal{G}[B]}$. It follows

that the restriction of the morphism $\mathcal{K} \rightarrow \overline{\mathcal{G}[B]} \cong \overline{\mathcal{K}[B]}$ to $\mathcal{K}[A]$ provides the required morphism. \square

The principal idea of the paper is to construct a series of E -generated permutation groups

$$G_1 \leftarrow G_2 \leftarrow \cdots \leftarrow G_{|E|} =: G \quad (3.4)$$

defined by an ascending sequence $\mathcal{X}_1 \subseteq \mathcal{X}_2 \subseteq \cdots \subseteq \mathcal{X}_{|E|}$ of complete E -graphs such that $G_k = \mathcal{T}(\mathcal{X}_k)$ is k -retractable and $G_{k+1} \rightarrow G_k$ is k -stable for every k . The crucial property of this sequence in relation to the given E -graph \mathcal{E} is the following:

For every word $p \in \tilde{E}^*$ on $k+1$ letters which forms a path $u \xrightarrow{p} v$ in \mathcal{E} and every letter $a \in A := \text{co}(p)$ either there is a word q in the letters $A \setminus \{a\}$ such that $[p]_{G_{k+1}} = [q]_{G_{k+1}}$ and q also forms a path $u \xrightarrow{q} v$ in \mathcal{E} , or otherwise (if no such q exists) there is a component in $\mathcal{X}_{k+1} \setminus \mathcal{X}_k$ which guarantees that G_{k+1} avoids the relation $p = p_{a \rightarrow 1}$, so that a belongs to the content of $[p]_{G_{k+1}[A]}$ and therefore to the content of $[p]_G$.

The graph-theoretic constructions to be introduced in the following are designed to serve this purpose. In order to guarantee that $G_{k+1} \rightarrow G_k$ is k -stable, the components \mathcal{L} of $\mathcal{X}_{k+1} \setminus \mathcal{X}_k$ will be assembled in a way that, for every k -element subset A of E , every A -component of \mathcal{L} appears as a subgraph of \mathcal{X}_k or is isomorphic with $\mathcal{G}_k[A]$. This turns out to be a challenging task.

3.3. Two crucial constructions: clusters and coset extensions. We introduce two crucial constructions involving Cayley graphs. Let G be an E -group; for $A \subseteq E$ and $g \in G$, $g\mathcal{G}[A]$ has the obvious meaning: it denotes the A -component of the vertex g of \mathcal{G} and is isomorphic (as an A -graph) with $\mathcal{G}[A]$ — we shall call such graphs *A -coset graphs* or simply *coset graphs* if the set of labels is understood. In the following subsections we shall construct new (bigger) graphs by gluing together disjoint copies of various coset graphs for different subsets $A \subseteq E$. In this context, the notation $v\mathcal{G}[A]$, where v is some vertex of a graph, means that the A -component of v in the graph in question is isomorphic with the full A -coset graph $\mathcal{G}[A]$.

Proviso 3.9. For the remainder of the section (§ 3.3.1–3) all E -groups G are assumed to be A -retractable, i.e. $G[A]$ is retractable for the (arbitrary but fixed) subset $A \subseteq E$ under consideration.

3.3.1. Clusters. Let G be an E -group, $A \subseteq E$ and assume that, as stated in Proviso 3.9, $G[A]$ is retractable. For every set \mathbb{P} of proper subsets of A , the graph

$$\text{CL}(G[A], \mathbb{P}) := \bigcup_{B \in \mathbb{P}} \mathcal{G}[B] \subseteq \mathcal{G}[A]$$

is an A -cluster. Note that $\text{CL}(G[A], \mathbb{P})$ is the subgraph of $\mathcal{G}[A]$ which is spanned by all B -paths in $\mathcal{G}[A]$ starting at 1, for $B \in \mathbb{P}$. The *core* of the cluster is the subgraph formed by the intersection $\bigcap_{B \in \mathbb{P}} \mathcal{G}[B]$, and by retractability of $G[A]$, $\bigcap_{B \in \mathbb{P}} \mathcal{G}[B] = \mathcal{G}[\bigcap_{B \in \mathbb{P}} B]$. This core is always non-empty but may consist of the vertex 1 only; the subgraphs $\mathcal{G}[B]$, for $B \in \mathbb{P}$, are the *constituent cosets* of the cluster $\text{CL}(G[A], \mathbb{P})$. Included in the definition of an A -cluster is, for $\mathbb{P} = \{B\}$, every graph $\mathcal{G}[B]$ with $B \subsetneq A$. The structure of $\text{CL}(G[A], \mathbb{P})$ as an A -graph actually only depends on the collection of the “small” subgroups $G[B]$, $B \in \mathbb{P}$ rather than on the entire group $G[A]$: indeed the cluster can be assembled from the constituents $\mathcal{G}[B]$ by forming their disjoint union and factoring by the congruence which identifies an element (vertex or edge) of some $\mathcal{G}[B]$ and some $\mathcal{G}[C]$ if and only if these elements coincide in $\mathcal{G}[B \cap C]$. More precisely, let $\varphi: \bigsqcup_{B \in \mathbb{P}} \mathcal{G}[B] \rightarrow \mathcal{G}$ be the morphism which maps every coset graph $\mathcal{G}[B]$ to itself, considered as a subgraph of \mathcal{G} . Let Θ be the mentioned congruence on $\bigsqcup_{B \in \mathbb{P}} \mathcal{G}[B]$. Then the kernel $\ker \varphi$ of φ (that is, the equivalence relation induced by φ on its domain) contains Θ ; retractability of $G[A]$ even implies the equality $\ker \varphi = \Theta$. From the Homomorphism Theorem we get

$$\text{CL}(G[A], \mathbb{P}) \cong \text{im}(\varphi) \cong \bigsqcup_{B \in \mathbb{P}} \mathcal{G}[B] / \Theta.$$

A consequence of this fact is the next lemma which will be of essential use in the proof of Proposition 5.4.

Lemma 3.10. *Let $G \twoheadrightarrow H$ be a $(k-1)$ -stable expansion between k -retractable E -groups G and H . Then for any $A \subseteq E$ with $|A| = k$ and any set \mathbb{P} of proper subsets of A , the labelled graphs $\text{CL}(G[A], \mathbb{P})$ and $\text{CL}(H[A], \mathbb{P})$ are isomorphic.*

Proof. This follows from the above discussion since $(k-1)$ -stability implies that $G[C] \cong H[C]$ for all $C \in \mathbb{P}$. \square

We next analyse the structure of B -components of A -clusters for $B \subsetneq A$. Let $\mathbb{P} = \{A_1, \dots, A_k\}$ be a set of proper subsets of A and let $B \subsetneq A$; let $v \in G[A]$ and $v\mathcal{G}[B]$ be the B -component of v in $\mathcal{G}[A]$. For the intersection of $v\mathcal{G}[B]$ with the cluster we have

$$\text{CL}(G[A], \mathbb{P}) \cap v\mathcal{G}[B] = \bigcup_{i=1}^k (\mathcal{G}[A_i] \cap v\mathcal{G}[B]).$$

The intersection $\mathcal{G}[A_i] \cap v\mathcal{G}[B]$ is either empty or a $(B \cap A_i)$ -coset $v_i\mathcal{G}[B \cap A_i]$ for some (any) $v_i \in \mathcal{G}[A_i] \cap v\mathcal{G}[B]$. For our purposes we may assume that $\mathcal{G}[A_i] \cap v\mathcal{G}[B] \neq \emptyset$ for every i . Indeed, we may assume that we have already removed those sets A_i for which $\mathcal{G}[A_i] \cap v\mathcal{G}[B] = \emptyset$.

Lemma 3.11. *If $\mathcal{G}[A_i] \cap v\mathcal{G}[B] \neq \emptyset$ for $i = 1, \dots, k$ then*

$$\mathcal{G}[A_1] \cap \dots \cap \mathcal{G}[A_k] \cap v\mathcal{G}[B] \neq \emptyset.$$

Proof. Let $t \leq k$ and assume that we have already proved that

$$\left(\bigcap_{i=1}^{t-1} \mathcal{G}[A_i] \right) \cap v\mathcal{G}[B] \neq \emptyset.$$

In order to simplify the notation we set $C := A_1 \cap \dots \cap A_{t-1}$ and $D := A_t$; then $\bigcap_{i=1}^{t-1} \mathcal{G}[A_i] = \mathcal{G}[C]$. So, let $u \in \left(\bigcap_{i=1}^{t-1} \mathcal{G}[A_i] \right) \cap v\mathcal{G}[B] = \mathcal{G}[C] \cap v\mathcal{G}[B]$ and $w \in \mathcal{G}[D] \cap v\mathcal{G}[B]$. Let $p \in \tilde{C}^*$ be such that $[p]_G = u^{-1}$ and $q \in \tilde{D}^*$




FIGURE 2

be such that $[q]_G = w^{-1}$; moreover, let $r \in \tilde{B}^*$ be a word which labels a path $u \rightarrow w$ running entirely in $v\mathcal{G}[B]$ (recall that all this happens in $\mathcal{G}[A]$). Let p_1 and q_1 be, respectively, the words obtained from p and q by deletion of all letters not in B . Since $[pq^{-1}]_G = [r]_G$ and $r \in \tilde{B}^*$, we have $[p_1q_1^{-1}]_G = [r]_G$, by retractability. Let $x := u \cdot p_1 = w \cdot q_1$. Then $p^{-1}p_1$ labels a path $1 \rightarrow x$ and so does $q^{-1}q_1$. Since $p^{-1}p_1 \in \tilde{C}^*$ and $q^{-1}q_1 \in \tilde{D}^*$, it follows that $x \in \mathcal{G}[C] \cap \mathcal{G}[D] = \mathcal{G}[C \cap D]$. From $x = u \cdot p_1$ and $p_1 \in \tilde{B}^*$ it follows that $x \in u\mathcal{G}[B] = v\mathcal{G}[B]$, altogether $x \in \mathcal{G}[B \cap D] \cap v\mathcal{G}[B]$. \square

The proof of Lemma 3.11 implicitly shows that retractable groups are 3-acyclic in the sense of condition (3.3), as stated in Remark 3.6. (Compare Figure 1 for coset patterns that are ruled out in the Cayley graph \mathcal{G} of any E -group G that is retractable; here now, the cosets in question, $1\mathcal{G}[C]$, $1\mathcal{G}[C]$ and $v\mathcal{G}[B]$, have x in their intersection, as indicated in Figure 2.)

In the situation of the proof of Lemma 3.11 we consider the automorphism of \mathcal{G} induced by left multiplication by x^{-1} for some $x \in G[A_1] \cap \dots \cap G[A_k] \cap vG[B]$, then $v\mathcal{G}[B] = x\mathcal{G}[B]$ and the intersection

$$\text{CL}(G[A], \mathbb{P}) \cap x\mathcal{G}[B] = \bigcup_{i=1}^k (\mathcal{G}[A_i] \cap x\mathcal{G}[B])$$

is isomorphic with the B -cluster $\text{CL}(G[B], \mathbb{O})$ where $\mathbb{O} = \{B \cap A_i : A_i \in \mathbb{P}\}$ (some of the sets $B \cap A_i$ may be empty, which perhaps degenerates to a full B -coset. This allows us to characterise the B -components of A -clusters for $B \subsetneq A$.

Corollary 3.12. *Let \mathbb{P} be a set of proper subsets of A and $B \subsetneq A$. Then every B -component of the cluster $\text{CL}(G[A], \mathbb{P})$ is either a B -coset, that is, isomorphic with $\mathcal{G}[B]$, or isomorphic with the B -cluster $\text{CL}(G[B], \mathbb{O})$ where $\mathbb{O} = \{C \cap B : C \in \mathbb{P}\}$ (some $C \cap B$ may be empty).*

Proof. The intersection $\text{CL}(G[A], \mathbb{P}) \cap v\mathcal{G}[B]$ is either the B -coset $v\mathcal{G}[B]$ itself (if it is contained in some constituent $\mathcal{G}[C]$ with $C \in \mathbb{P}$) or otherwise is isomorphic with the B -cluster $\text{CL}(G[B], \mathbb{O})$, as indicated above. Now let v be a vertex of $\text{CL}(G[A], \mathbb{P})$; then the B -component \mathcal{B} of v in $\text{CL}(G[A], \mathbb{P})$ is certainly contained in $\text{CL}(G[A], \mathbb{P}) \cap v\mathcal{G}[B]$. Since the latter intersection is a B -cluster, it is connected and therefore \mathcal{B} must coincide with this intersection. \square

Corollary 3.13. *Let $B, C \subsetneq A$; then the intersection $\mathcal{B} \cap \mathcal{C}$ of a B -component \mathcal{B} with a C -component \mathcal{C} of an A -cluster CL is either empty or a $B \cap C$ -coset or a $(B \cap C)$ -cluster.*

Proof. As mentioned above, $\mathcal{B} = \text{CL} \cap v\mathcal{G}[B]$ and $\mathcal{C} = \text{CL} \cap w\mathcal{G}[C]$ for some cosets $v\mathcal{G}[B]$ and $w\mathcal{G}[C]$. The latter two have either empty intersection or their intersection is a $(B \cap C)$ -coset $u\mathcal{G}[B \cap C]$ from which the claim follows. \square

We will need a generalisation of clusters, which we are going to present next. Let again be $G[A]$ be retractable (Proviso 3.9), \mathbb{P} be a set of proper subsets of A , v be a vertex of $\text{CL}(G[A], \mathbb{P})$ and $B \subsetneq A$. Under these assumptions we define

$$\text{CL}(G[A], \mathbb{P}) \circledast \mathcal{G}[B] := \bigcup_{C \in \mathbb{P}} \mathcal{G}[C] \cup v\mathcal{G}[B]$$

considered as a subgraph of $\mathcal{G}[A]$ and call the latter graph a *B -augmented A -cluster* or, more specifically, the *B -augmentation of $\text{CL}(G[A], \mathbb{P})$ at v* . We have already seen that the intersection $\text{CL}(G[A], \mathbb{P}) \cap v\mathcal{G}[B]$ is a B -component of $\text{CL}(G[A], \mathbb{P})$. It follows that the structure of the graph $\text{CL}(G[A], \mathbb{P}) \circledast \mathcal{G}[B]$ only depends on the collection $\{G[C] : C \in \mathbb{P}\}$, the vertex v and $G[B]$ rather than on the entire group $G[A]$. Indeed, as mentioned earlier, the structure of $\text{CL}(G[A], \mathbb{P})$ depends only on the graphs $\mathcal{G}[C]$ for $C \in \mathbb{P}$ and the B -component of v is a certain B -cluster \mathcal{B} which is isomorphic with a subgraph of $\mathcal{G}[B]$ via the monomorphism $\iota : \mathcal{B} \rightarrow \mathcal{G}[B]$ determined by $v \mapsto 1$. The augmented cluster $\text{CL}(G[A], \mathbb{P}) \circledast \mathcal{G}[B]$ then can be obtained as the disjoint union of $\text{CL}(G[A], \mathbb{P})$ and $\mathcal{G}[B]$ factored by the congruence whose non-singleton classes are $\{x, \iota(x)\}$ for all $x \in \mathcal{B}$ (x an edge or a vertex). As a consequence we obtain the following lemma, whose proof is analogous to the proof of Lemma 3.10; it will similarly be used in the proof of Proposition 5.4.

Lemma 3.14. *Let $G \twoheadrightarrow H$ be a $(k-1)$ -stable expansion between k -retractable E -groups G and H , φ the associated canonical morphism. Then, for any $A \subseteq E$ with $|A| = k$, any set \mathbb{P} of proper subsets of A , any $B \subsetneq A$ and any*

vertex u of $\text{CL}(G[A], \mathbb{P})$ with $v := \varphi(u)$ there is an isomorphism of labelled graphs

$$\text{CL}(G[A], \mathbb{P}) @ \mathcal{G}[B] \cong \text{CL}(H[A], \mathbb{P}) @ \mathcal{H}[B].$$

As the last result in this subsection we need to clarify, for $B, C \subsetneq A$, the structure of C -components of B -augmented A -clusters. These turn out to be $(B \cap C)$ -augmented C -clusters. We noticed already that every C -component of an A -cluster is a C -cluster (or a C -coset).

Corollary 3.15. *Let $B, C \subsetneq A$ and let $G[A]$ be retractable; then every C -component of a B -augmented A -cluster is a $(B \cap C)$ -augmented C -cluster (which includes C -clusters as a special case).*

Proof. Let the group G and $A, B, C \subseteq E$ be as in the statement of the corollary. Let $\text{CL}(G[A], \mathbb{P}) @ \mathcal{G}[B]$ be a B -augmentation of the A -cluster $\text{CL}(G[A], \mathbb{P})$ and let u be a vertex of this cluster. If the C -component \mathcal{C} of u in $\text{CL}(G[A], \mathbb{P})$ has empty intersection with the B -component \mathcal{B} of v in $\text{CL}(G[A], \mathbb{P})$ then \mathcal{C} coincides with the C -component of u in the augmented cluster and we are done as \mathcal{C} is a C -cluster (or a C -coset). Now assume that $\mathcal{C} \cap \mathcal{B} \neq \emptyset$ with w a vertex in $\mathcal{C} \cap \mathcal{B}$. We know that $\mathcal{C} \cap \mathcal{B}$ is a $(C \cap B)$ -cluster (Corollary 3.13) or a $(C \cap B)$ -coset and the C -component of w within $v\mathcal{G}[B] = w\mathcal{G}[B]$ consists exactly of the coset $w\mathcal{G}[B \cap C]$. It follows that the C -component of w in $\text{CL}(G[A], \mathbb{P}) @ \mathcal{G}[B]$ coincides with $\mathcal{C} \cup w\mathcal{G}[B \cap C] = \mathcal{C} @ \mathcal{G}[B \cap C]$ which is a $(B \cap C)$ -augmentation of the C -cluster \mathcal{C} . \square

3.3.2. Coset extensions. This second construction can be seen as a generalisation of clusters but is more involved. Let us fix an E -group G and a set $A \subseteq E$ of size $|A| \geq 2$. We assume that G is A -retractable, according to Proviso 3.9. Let \mathcal{K} be a connected A -subgraph of the Cayley graph \mathcal{G} of G . We recall that being an A -subgraph means that all labels of edges of \mathcal{K} belong to \tilde{A} (but not necessarily all such letters actually need to occur in \mathcal{K}). For some set $B \subsetneq A$ let $\mathcal{B} = v\mathcal{K}[B]$ be some B -component of \mathcal{K} ; this graph is embedded in $v\mathcal{G}[B] \cong \mathcal{G}[B]$. Moreover, for $B_1, B_2 \subsetneq B$ any B_1 - and B_2 -components \mathcal{B}_1 and \mathcal{B}_2 of \mathcal{B} are also embedded in $v\mathcal{G}[B]$ via their embedding in \mathcal{B} .

Definition 3.16 (admissibility for coset extension). Let G be an E -group, $A \subseteq E$ with $|A| \geq 2$, and assume that G is A -retractable (Proviso 3.9). Let \mathcal{K} be a connected A -subgraph of the Cayley graph \mathcal{G} of G . Consider all possible choices of subsets $B_1, B_2 \subsetneq B \subsetneq A$, of B -components $\mathcal{B} = v\mathcal{K}[B]$ of \mathcal{K} and for each pair of vertices $v_1, v_2 \in \mathcal{B}$ all possible B_1 - and B_2 -components $\mathcal{B}_1 = v_1\mathcal{B}[B_1] = v_1\mathcal{K}[B_1]$ and $\mathcal{B}_2 = v_2\mathcal{B}[B_2] = v_2\mathcal{K}[B_2]$. Then \mathcal{K} is *admissible for $\subseteq A$ -coset extension* (with respect to G) if

$$\mathcal{B}_1 \cap \mathcal{B}_2 = \emptyset \text{ in } \mathcal{B} \implies v_1\mathcal{G}[B_1] \cap v_2\mathcal{G}[B_2] = \emptyset \text{ in } v\mathcal{G}[B] \subseteq \mathcal{G}. \quad (3.5)$$

In other words, the patterns depicted in Figure 3 are forbidden in the context of a graph \mathcal{K} that is admissible for \mathcal{A} -coset extension (the right-hand side picture is for the case $B_1 = B_2$). The condition formulated in




FIGURE 3

Definition 3.16 corresponds to the notion of *freeness* in [19], here for the embedded graphs $\mathcal{B} = v\mathcal{K}[B]$ in $v\mathcal{G}[B]$. We note that, if \mathcal{K} is admissible for \mathcal{A} -coset extension, then, for every $B \subseteq A$, every B -component $v\mathcal{K}[B]$ is admissible for \mathcal{B} -coset extension.

Now let \mathcal{K} be a subgraph of \mathcal{G} that is admissible for \mathcal{A} -coset extension and fix a set $B \subsetneq A$. Let $\mathcal{B}_1, \dots, \mathcal{B}_k$ be all the B -components of \mathcal{K} . For every $i = 1, \dots, k$ select a vertex $v_i \in \mathcal{B}_i$. Then, in \mathcal{G} , the coset $v_i\mathcal{G}[B]$ contains \mathcal{B}_i as a subgraph. Let now $\text{CE}(G, \mathcal{K}; B)$ be the graph obtained by extending each component \mathcal{B}_i in \mathcal{K} to the entire coset $v_i\mathcal{G}[B]$. So $\text{CE}(G, \mathcal{K}; B)$ is the graph obtained by attaching in \mathcal{K} to each vertex v_i a copy $v_i\mathcal{G}[B]$ of $\mathcal{G}[B]$ and then identifying all of \mathcal{B}_i with its copy inside $v_i\mathcal{G}[B]$, but without performing any further identification (of vertices and/or edges). The graph $\text{CE}(G, \mathcal{K}; B)$ thus appears as a bunch of pairwise disjoint copies of $\mathcal{G}[B]$, connected by edges labelled by letters from $A \setminus B$. The union of the latter edges with all the \mathcal{B}_i then spans the graph \mathcal{K} .

We give a more formal definition of $\text{CE}(G, \mathcal{K}; B)$. Let \mathcal{K} be given with B -components $\mathcal{B}_1, \dots, \mathcal{B}_k$ and selected vertices $v_i \in \mathcal{B}_i$ for $i = 1, \dots, k$. For every i let $\iota_i: \mathcal{B}_i \rightarrow \mathcal{G}[B]$ be the unique graph monomorphism mapping v_i to 1. Then

$$\text{CE}(G, \mathcal{K}; B) := (\mathcal{K} \cup \bigcup_{i=1}^k \mathcal{G}[B] \times \{i\}) / \Theta \tag{3.6}$$

where Θ is the equivalence relation all of whose non-singleton equivalence classes are exactly the two-element sets

$$\{x, (\iota_i(x), i)\} \text{ with } x \in \mathcal{B}_i, i = 1, \dots, k$$

where x denotes a vertex or an edge of \mathcal{B}_i . The union on the right-hand side of (3.6) is a union of pairwise disjoint connected graphs and Θ is certainly

a congruence relation. The resulting graph $\text{CE}(G, \mathcal{K}; B)$ is the *B-coset extension of the A-graph* \mathcal{K} ; in this context, the subgraph \mathcal{K} of $\text{CE}(G, \mathcal{K}; B)$ is the *skeleton* of $\text{CE}(G, \mathcal{K}; B)$ and the subgraphs $v_i \mathcal{G}[B]$ are the *constituent cosets* of $\text{CE}(G, \mathcal{K}; B)$. For $C \subsetneq B \subsetneq A$, condition (3.5) of Definition 3.16 (by taking $B_1 = C = B_2$) implies that $\text{CE}(G, \mathcal{K}; C)$ is realised as a subgraph of $\text{CE}(G, \mathcal{K}; B)$. Moreover, for $C_1, C_2 \subsetneq B$, $C_1 \neq C_2$, once more condition (3.5) (this time taking $C_1 = B_1 \neq B_2 = C_2$) implies that

$$\text{CE}(G, \mathcal{K}; C_1) \cap \text{CE}(G, \mathcal{K}; C_2) = \text{CE}(G, \mathcal{K}; C_1 \cap C_2) \quad (3.7)$$

where the intersection takes place in $\text{CE}(G, \mathcal{K}; B)$. Now let \mathbb{P} be a set of proper subsets of A . Then the *\mathbb{P} -coset extension of \mathcal{K}* is defined as

$$\text{CE}(G, \mathcal{K}; \mathbb{P}) := \left(\bigcup_{B \in \mathbb{P}} \text{CE}(G, \mathcal{K}; B) \times \{B\} \right) / \Psi \quad (3.8)$$

where Ψ is the congruence defined on the disjoint union of all *B-coset extensions* $\text{CE}(G, \mathcal{K}; B)$ with $B \in \mathbb{P}$, by setting

$$(x_1, B_1) \Psi (x_2, B_2) :\iff x_1 = x_2 \in \text{CE}(G, \mathcal{K}; B_1 \cap B_2).$$

In other words, an edge or a vertex of $\text{CE}(G, \mathcal{K}; B_1)$ is identified with one in $\text{CE}(G, \mathcal{K}; B_2)$ if they represent the same element in $\text{CE}(G, \mathcal{K}; B_1 \cap B_2)$. Transitivity of Ψ follows from (3.7): indeed, for $i = 1, 2, 3$, let $B_i \in \mathbb{P}$ and $x_i \in \text{CE}(G, \mathcal{K}; B_i)$ be such that $(x_1, B_1) \Psi (x_2, B_2)$ and $(x_2, B_2) \Psi (x_3, B_3)$. Then

$$x_1 = x_2 \in \text{CE}(G, \mathcal{K}; B_1 \cap B_2) \text{ and } x_2 = x_3 \in \text{CE}(G, \mathcal{K}; B_2 \cap B_3)$$

so that

$$x_1 = x_3 \in \text{CE}(G, \mathcal{K}; B_1 \cap B_2) \cap \text{CE}(G, \mathcal{K}; B_2 \cap B_3) = \text{CE}(G, \mathcal{K}; B_1 \cap B_2 \cap B_3)$$

by application of (3.7) for $C_1 = B_1 \cap B_2$, $C_2 = B_2 \cap B_3$ and $B = B_2$, where the intersection takes place in $\text{CE}(G, \mathcal{K}; B_2)$. Provided that $B \in \mathbb{P}$, the coset extension $\text{CE}(G, \mathcal{K}; B)$ is embedded in $\text{CE}(G, \mathcal{K}; \mathbb{P})$ via $x \mapsto (x, B)\Psi$ where $(x, B)\Psi$ denotes the Ψ -class of (x, B) . For $v \in \mathcal{K}$ and $B \in \mathbb{P}$, the subgraphs $v \mathcal{G}[B]$ of $\text{CE}(G, \mathcal{K}; \mathbb{P})$ are the *constituent cosets* of $\text{CE}(G, \mathcal{K}; \mathbb{P})$ and the subgraph \mathcal{K} is the *skeleton* of $\text{CE}(G, \mathcal{K}; \mathbb{P})$.

Geometrically, the coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P})$ can be viewed as follows. For every $B \in \mathbb{P}$ consider $\text{CE}(G, \mathcal{K}; B)$ and attach these graphs to each other by identification of their skeleton \mathcal{K} , then form the largest *E-graph* quotient (that is, perform all identifications necessary to obtain an *E-graph*, but no more). The graph $\text{CE}(G, \mathcal{K}; \mathbb{P})$ then is the union

$$\text{CE}(G, \mathcal{K}; \mathbb{P}) = \bigcup_{B \in \mathbb{P}} \text{CE}(G, \mathcal{K}; B)$$

of its subgraphs $\text{CE}(G, \mathcal{K}; B)$ with $B \in \mathbb{P}$. For $B_1, B_2 \in \mathbb{P}$ then

$$\text{CE}(G, \mathcal{K}; B_1) \cap \text{CE}(G, \mathcal{K}; B_2) = \text{CE}(G, \mathcal{K}; B_1 \cap B_2). \quad (3.9)$$

This is reminiscent of (3.7) but B_1 and B_2 are now arbitrary members of \mathbb{P} (rather than subsets of some $B \subsetneq A$) and the intersection takes place in

$\text{CE}(G, \mathcal{K}; \mathbb{P})$ (rather than in $\text{CE}(G, \mathcal{K}; B)$). Moreover, condition (3.9) can be reformulated as a condition analogous to (3.5): for any $B_1, B_2 \in \mathbb{P}$ and vertices $v_1, v_2 \in \mathcal{K}$:

$$v_1\mathcal{K}[B_1] \cap v_2\mathcal{K}[B_2] = \emptyset \implies v_1\mathcal{G}[B_1] \cap v_2\mathcal{G}[B_2] = \emptyset \quad (3.10)$$

where the intersections take place in $\text{CE}(G, \mathcal{K}; \mathbb{P})$.

If every label of \mathcal{K} appears in some member B of \mathbb{P} , then $\text{CE}(G, \mathcal{K}; \mathbb{P})$ is weakly complete since every edge of $\text{CE}(G, \mathcal{K}; \mathbb{P})$ occurs in some coset subgraph $v\mathcal{G}[B]$. Most relevant will be the case $\mathbb{P} = \mathbb{P}_A$, the set of all proper subsets of A : we call $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ the *full* \mathcal{A} -coset extension of \mathcal{K} . In case $\mathcal{K} = \{v\}$ (one vertex, no edge) the \mathbb{P} -coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P})$ reduces to the cluster $\text{CL}(G[A], \mathbb{P})$.

Remark 3.17. An A -graph \mathcal{K} which is admissible for \mathcal{A} -coset extension may actually only contain edges labelled by letters (and their inverses) from some set $B \subsetneq A$. In this case $\text{CE}(G, \mathcal{K}; B) \cong \mathcal{G}[B]$; however, this is not in conflict with the definition of the full \mathcal{A} -coset extension. For sets $C \subsetneq A$ with $C \not\subseteq B$, the C -components of \mathcal{K} coincide with the $C \cap B$ -components, but nevertheless every such $C \cap B$ -component is extended to a full C -coset $v\mathcal{G}[C]$ in order to get $\text{CE}(G, \mathcal{K}; C)$.

We continue with further investigations of \mathcal{A} -coset extensions.

Proposition 3.18. *Let $\mathcal{K} \subseteq \mathcal{G}[A]$ be admissible for \mathcal{A} -coset extension and \mathbb{P} be a set of proper subsets of A . Then the inclusion monomorphism $\iota: \mathcal{K} \hookrightarrow \mathcal{G}[A]$ admits a unique extension to a graph morphism $\iota_{\mathbb{P}}: \text{CE}(G, \mathcal{K}; \mathbb{P}) \rightarrow \mathcal{G}[A]$.*

Proof. We first establish a unique extension $\iota_B: \text{CE}(G, \mathcal{K}; B) \rightarrow \mathcal{G}[A]$ for each $B \in \mathbb{P}$. Let $\mathcal{B}_1, \dots, \mathcal{B}_k$ be all B -components of \mathcal{K} with selected vertices $v_i \in \mathcal{B}_i$ for all i . Then for every i there is a unique graph monomorphism $\kappa_i: \mathcal{G}[B] \times \{i\} \rightarrow \mathcal{G}[A]$ such that $\kappa_i(1, i) = v_i$. The image of κ_i coincides with the coset subgraph $v_i\mathcal{G}[B]$ of $\mathcal{G}[A]$. Then, the union $\kappa := \iota \cup \bigcup_{i=1}^k \kappa_i$ is a morphism

$$\kappa: \mathcal{K} \cup \bigcup_{i=1}^k \mathcal{G}[B] \times \{i\} \rightarrow \mathcal{G}[A]$$

for which, for all i and $x \in \mathcal{B}_i$,

$$\kappa(x) = \iota(x) = x = \kappa_i(\iota_i(x), i) = \kappa(\iota_i(x), i)$$

where $\iota_i: \mathcal{B}_i \rightarrow \mathcal{G}[B]$ is the unique graph monomorphism mapping v_i to 1 that occurs in the definition of $\text{CE}(G, \mathcal{K}; B)$. It follows that the congruence Θ in (3.6) is contained in the kernel of κ and hence κ factors through $\text{CE}(G, \mathcal{K}; B)$ as $\kappa = \iota_B \circ \pi_{\Theta}$ (where π_{Θ} is the canonical projection $\pi_{\Theta}(x) = x\Theta$).

Next consider the disjoint union

$$\bigcup_{B \in \mathbb{P}} \text{CE}(G, \mathcal{K}; B) \times \{B\}$$

and let

$$\kappa_{\mathbb{P}} := \bigcup_{B \in \mathbb{P}} \iota_B: \bigcup_{B \in \mathbb{P}} \text{CE}(G, \mathcal{K}; B) \times \{B\} \rightarrow \mathcal{G}[A]$$

where $\iota_B: \text{CE}(G, \mathcal{K}; B) \times \{B\} \rightarrow \mathcal{G}[A]$ is defined by $\iota_B(x, B) = \iota_B(x)$. Similar to Θ and κ , the congruence Ψ that occurs in (3.8) is contained in the kernel of $\kappa_{\mathbb{P}}$, whence $\kappa_{\mathbb{P}}$ factors through $\text{CE}(G, \mathcal{K}; \mathbb{P})$ as $\kappa_{\mathbb{P}} = \iota_{\mathbb{P}} \circ \pi_{\Psi}$ for some unique morphism $\iota_{\mathbb{P}}: \text{CE}(G, \mathcal{K}; \mathbb{P}) \rightarrow \mathcal{G}[A]$ (with π_{Ψ} being again the projection $x \mapsto x\Psi$). \square

The morphism $\iota_B: \text{CE}(G, \mathcal{K}; B) \rightarrow \mathcal{G}[A]$ is injective when restricted either to the skeleton \mathcal{K} or to any constituent coset $v\mathcal{G}[B]$. However, in general ι_B is not injective on its entire domain $\text{CE}(G, \mathcal{K}; B)$. Within $\mathcal{G}[A]$ it may happen that for distinct vertices $v_i \neq v_j$ (as selected in the above proof) the corresponding cosets coincide: $v_i\mathcal{G}[B] = v_j\mathcal{G}[B] =: v\mathcal{G}[B]$. The coset $v\mathcal{G}[B]$ then contains (at least) two distinct B -components $\mathcal{B}_i \neq \mathcal{B}_j$ of \mathcal{K} . As a consequence, the vertices v_i and v_j can be connected by a B -path in $\mathcal{G}[A]$, but there is no B -path connecting these vertices in \mathcal{K} . This alludes to one of the key ideas of the paper and will eventually lead to the proof of the crucial Lemma 5.6.

Remark 3.19. Suppose that $H \twoheadrightarrow G$ is an expansion whose Cayley graph \mathcal{H} covers some completion of $\text{CE}(G, \mathcal{K}; B)$. Then the group H avoids every relation $p = q$ where p is any word labelling a path in \mathcal{K} that connects two distinct B -components of \mathcal{K} and q is any \tilde{B} -word, essentially because the graph $\text{CE}(G, \mathcal{K}; B)$ unfolds the subgraph $\mathcal{K} \cup \bigcup v_i\mathcal{G}[B]$ of $\mathcal{G}[A]$ that arises as the image of $\text{CE}(G, \mathcal{K}; B)$ under ι_B (see Figure 4).




FIGURE 4. Part of $\mathcal{K} \cup \bigcup_{t=1}^k v_t\mathcal{G}[B] \subseteq \mathcal{G}[A]$ and of $\text{CE}(G, \mathcal{K}; B)$

We are able to sharpen an earlier remark. Let \mathcal{K} be a connected A -graph admissible for \mathcal{A} -coset extension, let $B \subsetneq A$ and let $\mathcal{B} \subseteq \mathcal{K}$ be a B -component of \mathcal{K} with v a vertex of \mathcal{B} . By construction of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$,

$$v \in \mathcal{B} \subseteq v\mathcal{G}[B] \subseteq \text{CE}(G, \mathcal{K}; \mathbb{P}_A).$$

But \mathcal{B} is itself admissible for \mathcal{B} -coset extension and hence $\text{CE}(G, \mathcal{B}; \mathbb{P}_B)$ is well defined. Admissibility of \mathcal{K} (Definition 3.16) implies that in this case

the morphism $\iota_{\mathbb{P}_B} : \text{CE}(G, \mathcal{B}; \mathbb{P}_B) \rightarrow \mathcal{G}[B]$ of Proposition 3.18 is injective, so that we get the following.

Lemma 3.20. *Let \mathcal{K} be a subgraph of $\mathcal{G}[A]$ which is admissible for \mathcal{A} -coset extension (in particular $G[A]$ is retractable, cf. Definition 3.16 and also Proviso 3.9). Let $B \subsetneq A$ with $|B| \geq 2$; then every B -component \mathcal{B} of \mathcal{K} is admissible for \mathcal{B} -coset extension and the morphism $\iota_{\mathbb{P}_B} : \text{CE}(G, \mathcal{B}; \mathbb{P}_B) \rightarrow \mathcal{G}[B]$ is injective. In particular, for any vertex $v \in \mathcal{B}$,*

$$v \in \mathcal{B} \subseteq \text{CE}(G, \mathcal{B}; \mathbb{P}_B) \subseteq v\mathcal{G}[B] \subseteq \text{CE}(G, \mathcal{K}; \mathbb{P}_A).$$

Another consequence concerns connectivity in the graph \mathcal{K} ; it will be of significant use later. In terms of [19] this means that a graph \mathcal{K} which is admissible for \mathcal{A} -coset extension is 2-acyclic.

Lemma 3.21. *Suppose that the graph $\mathcal{K} \subseteq \mathcal{G}$ is admissible for \mathcal{A} -coset extension. Then, for any $B, C \subsetneq A$, the intersection $\mathcal{B} \cap \mathcal{C}$ of any B -component \mathcal{B} and any C -component \mathcal{C} of \mathcal{K} is connected and hence is a $(B \cap C)$ -component.*

Proof. Suppose that $B \neq C$ and let u, v be vertices of $\mathcal{B} \cap \mathcal{C}$ and assume that they belong to distinct components of $\mathcal{B} \cap \mathcal{C}$. Admissibility of \mathcal{K} (by taking $B_1 = B \cap C = B_2$) implies that the cosets $u\mathcal{G}[B \cap C]$ and $v\mathcal{G}[B \cap C]$ are disjoint (that is, distinct), and both cosets are contained in $u\mathcal{G}[B] = v\mathcal{G}[B]$ as well as $u\mathcal{G}[C] = v\mathcal{G}[C]$. Consider the graph morphism $\iota_{\mathbb{P}_A} : \text{CE}(G, \mathcal{K}; \mathbb{P}_A) \rightarrow \mathcal{G}[A]$. It maps the cosets $u\mathcal{G}[B]$ as well as $v\mathcal{G}[C]$ injectively to the corresponding coset subgraphs of $\mathcal{G}[A]$. Since $u\mathcal{G}[B \cap C]$ and $v\mathcal{G}[B \cap C]$ are disjoint, it follows that the intersection of the cosets $u\mathcal{G}[B]$ and $v\mathcal{G}[C]$ (in $\mathcal{G}[A]$) is disconnected as it has at least the two components $u\mathcal{G}[B \cap C]$ and $v\mathcal{G}[B \cap C]$; this, however, contradicts the assumption that $G[A]$ is retractable. \square

3.3.3. Augmented coset extensions. Similarly to augmented clusters we require augmented coset extensions. Again fix an E -group G , let $A \subseteq E$ with $|A| \geq 2$ and assume that $G[A]$ is retractable, according to Proviso 3.9. Let $\mathcal{K} \subseteq \mathcal{G}[A]$ be admissible for \mathcal{A} -coset extension. Recall that the full \mathcal{A} -coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ can be seen as the union $\bigcup_{B \subsetneq A} \text{CE}(G, \mathcal{K}, B)$ where for $B, C \subsetneq A$,

$$\text{CE}(G, \mathcal{K}; B) \cap \text{CE}(G, \mathcal{K}; C) = \text{CE}(G, \mathcal{K}, B \cap C).$$

Every vertex x of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ is sitting in some $\text{CE}(G, \mathcal{K}; B)$, and, inside $\text{CE}(G, \mathcal{K}; B)$ in a unique *constituent coset* $v\mathcal{G}[B]$ with $v \in \mathcal{K}$. The vertex v is not unique, but unique is its B -component $v\mathcal{K}[B]$. In this situation we say that the pair (B, v) *supports* the vertex x or *provides support* for the vertex x in $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$; the *size* of this support is $|B|$. This actually means that the skeleton \mathcal{K} may be accessed from the vertex x by a B -path whose terminal vertex is v . We say that (B, v) provides *unique minimal support* if, whenever (C, w) provides support for x then $B \subseteq C$ and $v\mathcal{K}[B] \subseteq w\mathcal{K}[C]$. Now let \mathcal{J} be a subgraph of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$; for a set $B \subsetneq A$ and a vertex

$v \in \mathcal{K}$ we say that (B, v) *provides unique minimal support for* \mathcal{J} , or that \mathcal{J} *has unique minimal support through* (B, v) , if (B, v) supports some vertex x of \mathcal{J} , and if some pair (C, w) supports any vertex y of \mathcal{J} then $B \subseteq C$ and $v\mathcal{K}[B] \subseteq w\mathcal{K}[C]$. In this case we say that the unique minimal support of \mathcal{J} is *attained* at the vertex x . Notice that the condition $v\mathcal{K}[B] \subseteq w\mathcal{K}[C]$ implies the inclusion $v\mathcal{G}[B] \subseteq w\mathcal{G}[C] = v\mathcal{G}[C]$ for the constituent cosets involved. From (3.9) it follows that every one-vertex subgraph of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ has unique minimal support.

We come to a crucial property, which the full \mathcal{A} -coset extension of a graph \mathcal{K} may or may not have.

Definition 3.22 (cluster property). The full coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ has the *cluster property* if, for every $B \subsetneq A$ the following hold:

- (1) every B -component \mathcal{B} of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ which has empty intersection with the skeleton \mathcal{K} is a B -cluster or a full B -coset;
- (2) every \mathcal{B} of (1) has unique minimal support which is attained at some vertex x of the core of \mathcal{B} (if \mathcal{B} is a cluster).

Note that minimal support will typically not be attained at all core vertices. We first show that the cluster property implies that components of the coset extension intersect nicely, that is, the coset extension is 2-acyclic in terms of [19].

Proposition 3.23. *Suppose that $\mathcal{K} \subseteq \mathcal{G}[A]$ is admissible for \mathcal{A} -coset extension and that the full \mathcal{A} -coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ has the cluster property. Then, for all pairs $B, C \subsetneq A$ the intersection $\mathcal{B} \cap \mathcal{C}$ of any B -component \mathcal{B} and any C -component \mathcal{C} is connected and hence is a $(B \cap C)$ -component of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$.*

Proof. We consider several cases and start with the most difficult one: suppose that both \mathcal{B} and \mathcal{C} have empty intersection with the skeleton \mathcal{K} . We need to show that $\mathcal{B} \cap \mathcal{C}$ is connected. We know that \mathcal{B} is a B -cluster, \mathcal{C} is a C -cluster, that is, $\mathcal{B} \cong \text{CL}(G[B], \{B_1, \dots, B_k\})$ and $\mathcal{C} \cong \text{CL}(G[C], \{C_1, \dots, C_l\})$ for $B_i \subsetneq B$ and $C_j \subsetneq C$; it may also happen that $k = 1$ and/or $l = 1$ in which case it may happen that $B_1 = B$ and/or $C_1 = C$ (that is, \mathcal{B} and/or \mathcal{C} is a B -coset and/or C -coset) — the argument for this subcase is similar but simpler. Let x be a vertex in the core of \mathcal{B} , y a vertex in the core of \mathcal{C} , such that the unique minimal support (M, m) of \mathcal{B} is attained at x , and the unique minimal support (N, n) of \mathcal{C} is attained at y . Then $\mathcal{B} = \bigcup_{i=1}^k x\mathcal{G}[B_i]$ and $\mathcal{C} = \bigcup_{j=1}^l y\mathcal{G}[C_j]$. Let $u_1 \neq u_2$ be vertices of $\mathcal{B} \cap \mathcal{C}$; we may assume that $u_1 \in x\mathcal{G}[B_1] \cap y\mathcal{G}[C_1]$ and $u_2 \in x\mathcal{G}[B_2] \cap y\mathcal{G}[C_2]$. The vertices u_1 and u_2 also have unique minimal support (F_1, v_1) and (F_2, v_2) , say. Then $M, N \subseteq F_1, F_2$ and even more holds, namely

$$\begin{aligned} m\mathcal{G}[M], n\mathcal{G}[N] &\subseteq m\mathcal{G}[F_1] = v_1\mathcal{G}[F_1] = n\mathcal{G}[F_1] \\ \text{and } m\mathcal{G}[M], n\mathcal{G}[N] &\subseteq m\mathcal{G}[F_2] = v_2\mathcal{G}[F_2] = n\mathcal{G}[F_2]. \end{aligned}$$

The equality $m\mathcal{G}[F_1] = v_1\mathcal{G}[F_1]$ follows from the fact that (F_1, v_1) provides some support for \mathcal{B} , while (M, m) provides unique minimal support for \mathcal{B} hence $M \subseteq F_1$ and $m \in m\mathcal{G}[M] \subseteq v_1\mathcal{G}[F_1]$; likewise, (F_1, v_1) provides some support for \mathcal{C} while (N, n) provides unique minimal support for \mathcal{C} , hence $N \subseteq F_1$ and $n \in n\mathcal{G}[N] \subseteq v_1\mathcal{G}[F_1]$ which implies $v_1\mathcal{G}[F_1] = n\mathcal{G}[F_1]$. The remaining two equalities are proved in the same fashion. From

$$m\mathcal{G}[M] \cup n\mathcal{G}[N] \subseteq v_1\mathcal{G}[F_1] \cap v_2\mathcal{G}[F_2]$$

we get $v_1\mathcal{G}[F_1] \cap v_2\mathcal{G}[F_2] \neq \emptyset$, which by (3.10) implies $v_1\mathcal{K}[F_1] \cap v_2\mathcal{K}[F_2] \neq \emptyset$. By Lemma 3.21, this intersection is an F -component of \mathcal{K} for $F = F_1 \cap F_2$, that is,

$$v_1\mathcal{K}[F_1] \cap v_2\mathcal{K}[F_2] = m\mathcal{K}[F] = n\mathcal{K}[F].$$

From the definition of the full coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ and (3.9) it follows that the intersection $v_1\mathcal{G}[F_1] \cap v_2\mathcal{G}[F_2]$ itself is connected (it is isomorphic with $m\mathcal{G}[F] = n\mathcal{G}[F]$). So the subgraph of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ formed by the union $v_1\mathcal{G}[F_1] \cup v_2\mathcal{G}[F_2]$ is isomorphic with the cluster $\text{CL}(G[A], \{F_1, F_2\})$, see Figure 5.

Moreover, the cosets $x\mathcal{G}[B_1]$ and $v_1\mathcal{G}[F_1]$ both are contained in some constituent coset $w\mathcal{G}[D]$. Indeed, $x\mathcal{G}[B_1]$ arises as the intersection of the B -component \mathcal{B} with some constituent coset, say $w\mathcal{G}[D]$, for some vertex $w \in \mathcal{K}$ and $D \subsetneq A$. Then (D, w) supports u_1 , whence $F_1 \subseteq D$ and $v_1\mathcal{G}[F_1] \subseteq v_1\mathcal{G}[D] = w\mathcal{G}[D]$. Since $G[D]$ is retractable the intersection $x\mathcal{G}[B_1] \cap v_1\mathcal{G}[F_1]$ is connected. The same holds for the intersections

$$x\mathcal{G}[B_2] \cap v_2\mathcal{G}[F_2], \quad y\mathcal{G}[C_1] \cap v_1\mathcal{G}[F_1] \quad \text{and} \quad y\mathcal{G}[C_2] \cap v_2\mathcal{G}[F_2].$$

Setting $B' := (B_1 \cap F_1) \cup (B_2 \cap F_2)$ and $C' := (C_1 \cap F_1) \cup (C_2 \cap F_2)$ we see that u_1 and u_2 belong to the same B' - as well as C' -component of the cluster $v_1\mathcal{G}[F_1] \cup v_2\mathcal{G}[F_2]$, the intersection of which is a $(B' \cap C')$ -component of that cluster, by Corollary 3.13. Consequently, u_1 and u_2 are in the same $(B' \cap C')$ -component of $v_1\mathcal{G}[F_1] \cup v_2\mathcal{G}[F_2]$ and hence in the same $(B \cap C)$ -component of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$; the configuration is depicted in Figure 5.

Next we consider the case when \mathcal{C} has empty intersection with the skeleton \mathcal{K} (as in the previous case), but \mathcal{B} has not. Then $\mathcal{C} \cong \text{CL}(G[C], \{C_1, \dots, C_l\})$ and $\mathcal{B} = v\mathcal{G}[B]$ for some vertex $v \in \mathcal{K}$. We let $u_1 \neq u_2$ be vertices in $\mathcal{B} \cap \mathcal{C}$, and we may assume that $u_1 \in \mathcal{C}_1 := y\mathcal{G}[C_1]$ and $u_2 \in \mathcal{C}_2 := y\mathcal{G}[C_2]$ (as in the previous case), where y is a vertex in the core of \mathcal{C} which attains minimal support of \mathcal{C} . In this case (B, v) supports u_1 as well as u_2 and therefore also y , so that $u_1, y, u_2 \in \mathcal{B} = v\mathcal{G}[B]$, see Figure 6. For the same reason as in the previous case, the intersections $y\mathcal{G}[C_1] \cap v\mathcal{G}[B]$ and $y\mathcal{G}[C_2] \cap v\mathcal{G}[B]$ both are connected. Hence there is a $(B \cap C)$ -path $u_1 \rightarrow y$ and also a $(B \cap C)$ -path $y \rightarrow u_2$, and altogether there is a $(B \cap C)$ -path $u_1 \rightarrow u_2$.

Finally, the case when \mathcal{B} as well as \mathcal{C} have non-empty intersection with the skeleton \mathcal{K} is obvious, since in this case $\mathcal{B} \cap \mathcal{C}$ is a $(B \cap C)$ -coset. \square

We are led to a further construction. Let \mathcal{K} be admissible for \mathcal{A} -coset extension and suppose that the full \mathcal{A} -coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ has the

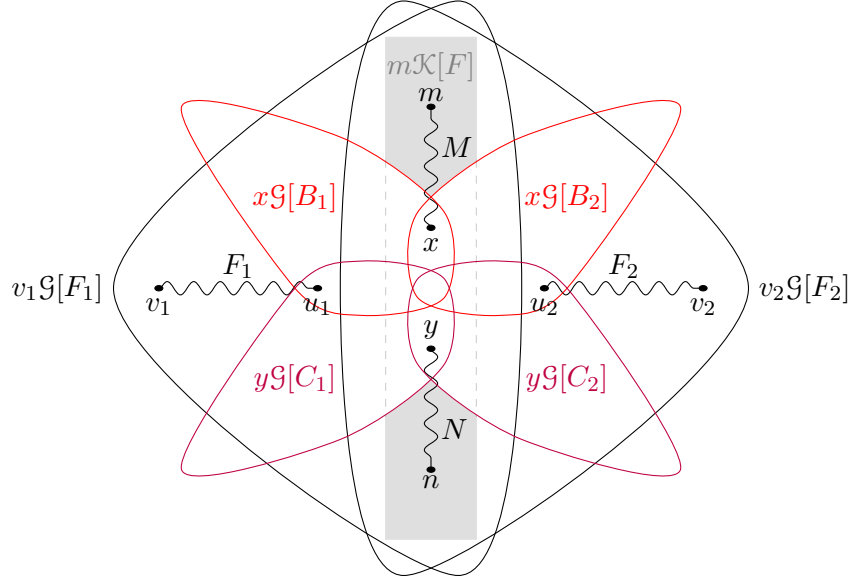


FIGURE 5




FIGURE 6

cluster property. For a vertex $v \in \text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ and some $B \subsetneq A$ the B -component \mathcal{B} of v is either a B -coset $v\mathcal{G}[B]$ (in this case, \mathcal{B} may or may not intersect with the skeleton \mathcal{K}) or a proper B -cluster (in which case it does not intersect with the skeleton \mathcal{K}). In any case, \mathcal{B} embeds into $\mathcal{G}[B]$ via some graph monomorphism $\iota: \mathcal{B} \hookrightarrow \mathcal{G}[B]$ (which is unique if one additionally assumes that $\iota(v) = 1$). We define the B -augmentation at v of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ by

$$\text{CE}(G, \mathcal{K}; \mathbb{P}_A) \odot \mathcal{G}[B] := \text{CE}(G, \mathcal{K}; \mathbb{P}_A) \sqcup \mathcal{G}[B] / \Omega$$

where Ω is the congruence whose non-singleton congruence classes are the two-element sets $\{x, \iota(x)\}$ for $x \in \mathcal{B}$. We note that $\text{CE}(G, \mathcal{K}; \mathbb{P}_A) \odot \mathcal{G}[B]$ can be written as the union

$$\text{CE}(G, \mathcal{K}; \mathbb{P}_A) \cup v\mathcal{G}[B]$$

of its two subgraphs $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ and $v\mathcal{G}[B]$ whose intersection is just the B -component \mathcal{B} of v in $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$.

Proposition 3.24. *Let $B, C \subsetneq A$ and \mathcal{K} be admissible for $\sphericalangle A$ -coset extension and such that the full $\sphericalangle A$ -coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ enjoys the cluster property. Then every C -component of any B -augmented full coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P}_A) \textcircled{+} \mathcal{G}[B]$ is either a C -coset, a $B \cap C$ -coset, a C -cluster or a $(B \cap C)$ -augmented C -cluster.*

Proof. Let \mathcal{C} be a C -component of $\text{CE}(G, \mathcal{K}; \mathbb{P}_A) \textcircled{+} \mathcal{G}[B]$. If $\mathcal{C} \subseteq \text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ or $\mathcal{C} \subseteq v\mathcal{G}[B]$ we are done: \mathcal{C} happens to be a C -coset or a $B \cap C$ -coset or a C -cluster. Let us assume that \mathcal{C} is contained neither in $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ nor in $v\mathcal{G}[B]$. We have

$$\mathcal{C} = \underbrace{(\text{CE}(G, \mathcal{K}; \mathbb{P}_A) \cap \mathcal{C})}_{\mathcal{C}_1} \cup \underbrace{(v\mathcal{G}[B] \cap \mathcal{C})}_{\mathcal{C}_2}$$

and \mathcal{C}_1 is a proper C -cluster (if it were a C -coset it would coincide with \mathcal{C} , which would be contained in $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$). Let \mathcal{B}_v be the B -component of v in $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$. Our assumption implies that $\mathcal{C} \cap \mathcal{B}_v \neq \emptyset$. Let w be a vertex of $\mathcal{C} \cap \mathcal{B}_v$. By Proposition 3.23, $\mathcal{C} \cap \mathcal{B}_v = \mathcal{C}_1 \cap \mathcal{B}_v$ is the $(B \cap C)$ -component of w in $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$, which is a $(B \cap C)$ -cluster or a $(B \cap C)$ -coset. Moreover,

$$\mathcal{C}_2 = \mathcal{C} \cap v\mathcal{G}[B] = \mathcal{C} \cap w\mathcal{G}[B] = w\mathcal{G}[B \cap C].$$

If $\mathcal{C}_1 \cap \mathcal{B}_v$ were a $(B \cap C)$ -coset, then it would coincide with $w\mathcal{G}[B \cap C]$ and again $\mathcal{C} \subseteq \text{CE}(G, \mathcal{K}; \mathbb{P}_A)$. Hence, under our assumption, $\mathcal{C}_1 \cap \mathcal{B}_v$ is indeed a proper $(B \cap C)$ -cluster. So we see that $\mathcal{C} = \mathcal{C}_1 \cup w\mathcal{G}[B \cap C]$ and

$$\mathcal{C}_1 \cap w\mathcal{G}[B \cap C] = \text{CE}(G, \mathcal{K}; \mathbb{P}_A) \cap w\mathcal{G}[B \cap C]$$

is the $(B \cap C)$ -component of w in $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$. Altogether this just means that $\mathcal{C} = \mathcal{C}_1 \textcircled{+} \mathcal{G}[B \cap C]$, that is, \mathcal{C} is the $(B \cap C)$ -augmentation of the C -cluster \mathcal{C}_1 at w . \square

4. THE TWO MAIN TECHNICAL RESULTS

In this section we prove the two main technical results. They will be essential to set up the inductive procedure to gain the series (3.4). In order to do so, we need another crucial definition. Assume, as above, that $|A| \geq 2$, that $G[A]$ is retractable and that $\mathcal{K} \subseteq \mathcal{G}[A]$ is admissible for $\sphericalangle A$ -coset extension.

Definition 4.1 (bridge freeness). The full coset extension $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$ is *bridge-free* in $\mathcal{G}[A]$ if

- (1) the morphism $\iota_{\mathbb{P}_A} : \text{CE}(G, \mathcal{K}; \mathbb{P}_A) \rightarrow \mathcal{G}[A]$ (Proposition 3.18) is an embedding;
- (2) for every $B \subsetneq A$, if two vertices $u, v \in \text{CE}(G, \mathcal{K}; \mathbb{P}_A) \subseteq \mathcal{G}[A]$ (as per (1)) are B -connected in $\mathcal{G}[A]$ then they are B -connected even in $\text{CE}(G, \mathcal{K}; \mathbb{P}_A)$.

In Theorem 4.2 take into account Lemma 3.20: if some subgraph $\mathcal{L} \subseteq \mathcal{H}[A]$ of the Cayley graph of the group H is admissible for $\mathcal{C}A$ -coset extension, then all its B -components $v\mathcal{L}[B]$, for $B \subsetneq A$, are admissible for $\mathcal{C}B$ -coset extension and the morphisms of Proposition 3.18 are embeddings $\text{CE}(H, v\mathcal{L}[B]; \mathbb{P}_B) \hookrightarrow v\mathcal{H}[B]$.

Theorem 4.2 (forward induction). *Let H be an E -group, $A \subseteq E$, $|A| \geq 3$ and suppose that $H[A]$ is retractable. Let $\mathcal{L} \subseteq \mathcal{H}[A]$ be a connected A -graph which is admissible for $\mathcal{C}A$ -coset extension. Assume that for all $B \subsetneq A$ and every vertex $v \in \mathcal{L}$, the full $\mathcal{C}B$ -coset extension $\text{CE}(H, v\mathcal{L}[B]; \mathbb{P}_B)$*

- (1) *is bridge-free in $\mathcal{H}[B]$, and*
- (2) *has the cluster property.*

Then the full $\mathcal{C}A$ -coset extension $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$ has the cluster property.

Proof. Let $B \subsetneq A$, let \mathcal{B} be a B -component of $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$ and suppose that \mathcal{B} has empty intersection with the skeleton \mathcal{L} . We first show the following: if \mathcal{B} is not fully contained in any one constituent coset of $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$, then the intersection of \mathcal{B} with any constituent coset is either empty or contains a vertex that is supported by fewer than $|A| - 1$ elements. Indeed, let $\mathcal{B} \cap v_1\mathcal{H}[A_1] \neq \emptyset$, w.l.o.g. $|A_1| = |A| - 1$, and assume that \mathcal{B} is not contained in $v_1\mathcal{H}[A_1]$. Then some vertex $s_1 \in \mathcal{B} \cap v_1\mathcal{H}[A_1]$ must be connected by an edge e in \mathcal{B} to some vertex $s_2 \in (\mathcal{B} \cap v_2\mathcal{H}[A_2]) \setminus v_1\mathcal{H}[A_1]$ in some other constituent coset $v_2\mathcal{H}[A_2]$, for some $A_2 \neq A_1$. Since $s_2 \notin v_1\mathcal{H}[A_1]$, also $e \notin v_1\mathcal{H}[A_1]$. Then e belongs to a coset $v_3\mathcal{H}[A_3]$ (possibly coinciding with $v_2\mathcal{H}[A_2]$) with $A_3 \neq A_1$. In any case, $s_1, s_2 \in v_3\mathcal{H}[A_3]$ (if a graph contains an edge then also its initial and terminal vertices). It follows that s_1 is supported by (A_3, v_3) , that is, $s_1 \in v_1\mathcal{H}[A_1] \cap v_3\mathcal{H}[A_3] = v\mathcal{H}[A_1 \cap A_3]$ for some vertex v , and $|A_1 \cap A_3| < |A| - 1$.

Therefore, if no vertex of \mathcal{B} has support of size smaller than $|A| - 1$, then \mathcal{B} is contained in some constituent coset $v_1\mathcal{H}[A_1]$ with $|A_1| = |A| - 1$, and therefore is a $B \cap A_1$ -coset with minimal support (A_1, v_1) .

We are left with the case that \mathcal{B} admits support of size strictly smaller than $|A| - 1$. We collect some constituent cosets $v_1\mathcal{H}[A_1], \dots, v_n\mathcal{H}[A_n]$ of $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$ for generator sets $A_i \subsetneq A$ of size $|A_i| = |A| - 1$ such that $\mathcal{B} \subseteq \bigcup_{i=1}^n v_i\mathcal{H}[A_i]$ and we assume that the choice of the constituent cosets $v_i\mathcal{H}[A_i]$ is minimal for $\mathcal{B} \subseteq \bigcup_{i=1}^n v_i\mathcal{H}[A_i]$ in the sense that \mathcal{B} is not contained in any union of fewer than n constituent cosets. Then

$$\mathcal{B} = \mathcal{B} \cap \left(\bigcup_{i=1}^n v_i\mathcal{H}[A_i] \right) = \bigcup_{i=1}^n (\mathcal{B} \cap v_i\mathcal{H}[A_i]) = \bigcup_{i=1}^n \mathcal{B}_i$$

for $\mathcal{B}_i = \mathcal{B} \cap v_i\mathcal{H}[A_i]$. Every \mathcal{B}_i is a non-empty B_i -coset subgraph of $v_i\mathcal{H}[A_i]$ where $B_i = B \cap A_i$ and all B_i have size at most $|A| - 2$. (If for some i , $|B_i| = |A| - 1$ then $B_i = A_i$ and $\mathcal{B}_i = v_i\mathcal{H}[A_i]$ would have non-empty intersection with the skeleton \mathcal{L} .) In addition, every \mathcal{B}_i has a vertex supported by fewer than $|A_i| = |A| - 1$ letters: if $n = 1$ this is immediate and if $n > 1$ then \mathcal{B} is

not contained in a single constituent coset, and the situation is as discussed at the start of the proof.

We need to verify items (1) and (2) of Definition 3.22. For $i = 1, \dots, n$ denote by \mathcal{A}_i the A_i -component $v_i\mathcal{L}[A_i]$ of v_i in \mathcal{L} . By Lemma 3.20, \mathcal{A}_i is admissible for $\sphericalangle A_i$ -coset extension and the full $\sphericalangle A_i$ -coset extension $\text{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i})$ embeds into $v_i\mathcal{H}[A_i]$ (via the mapping of Proposition 3.18). Since \mathcal{B}_i admits vertices supported by fewer than $|A_i| = |A| - 1$ letters, we have that $\mathcal{B}_i \cap \text{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i}) \neq \emptyset$ — once more we take into account that

$$\text{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i}) \subseteq v_i\mathcal{H}[A_i] \subseteq \text{CE}(H, \mathcal{L}; \mathbb{P}_A).$$

Bridge-freeness of $\text{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i})$ (assumption (1)) implies that

$$\mathcal{B}_i \cap \text{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i})$$

is connected. By assumption (2) therefore, $\mathcal{B}_i \cap \text{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i})$ has unique minimal support in $\text{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i})$, say (D_i, u_i) . But then the pair (D_i, u_i) also provides unique minimal support of \mathcal{B}_i in $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$. If $n = 1$ we are already done; so let us assume that $n \geq 2$. Minimality of (D_i, u_i) implies in particular that any path connecting \mathcal{B}_i to \mathcal{L} in $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$ must use (at least) all labels in D_i and, in case it uses only labels from D_i , necessarily leads to the D_i -component $u_i\mathcal{A}_i[D_i]$. So, for every i , there exist vertices $s_i \in \mathcal{B}_i$, $u_i \in \mathcal{A}_i$ and a word $m_i \in \widetilde{D}_i^*$ labelling a path $s_i \rightarrow u_i$ which runs entirely inside the coset $u_i\mathcal{H}[D_i]$, which in turn is contained in $v_i\mathcal{H}[A_i] = u_i\mathcal{H}[A_i]$.

Since $\mathcal{B} = \bigcup_{i=1}^n \mathcal{B}_i$ is connected, there are i, j such that $\mathcal{B}_i \cap \mathcal{B}_j \neq \emptyset$; after some renumbering we may assume that $\mathcal{B}_1 \cap \mathcal{B}_2 \neq \emptyset$. Then also $v_1\mathcal{H}[A_1] \cap v_2\mathcal{H}[A_2] \neq \emptyset$; from (3.10) we get $\mathcal{A}_1 \cap \mathcal{A}_2 \neq \emptyset$ and by Lemma 3.21, $\mathcal{A}_1 \cap \mathcal{A}_2$ is an $(A_1 \cap A_2)$ -component of \mathcal{L} , say $v\mathcal{L}[A_1 \cap A_2]$ for some $v \in \mathcal{A}_1 \cap \mathcal{A}_2$. From (3.9) it follows that $v_1\mathcal{H}[A_1] \cap v_2\mathcal{H}[A_2] = v\mathcal{H}[A_1 \cap A_2]$. The intersection $\mathcal{B}_1 \cap \mathcal{B}_2$ is a $B \cap A_1 \cap A_2$ coset in $v\mathcal{H}[A_1 \cap A_2]$. Similarly as for \mathcal{B}_1 one argues that $\mathcal{B}_1 \cap \mathcal{B}_2$ has unique minimal support in $\text{CE}(H, \mathcal{A}_1 \cap \mathcal{A}_2; \mathbb{P}_{A_1 \cap A_2})$, (D, u) say, which (as for \mathcal{B}_1) provides unique minimal support of $\mathcal{B}_1 \cap \mathcal{B}_2$ in $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$. Let $s \in \mathcal{B}_1 \cap \mathcal{B}_2$ be a vertex which attains the support (D, u) . So far, the situation is depicted as in Figure 7. We note that $D \subseteq A_1 \cap A_2$




FIGURE 7

and so

$$u\mathcal{H}[D] \subseteq u\mathcal{H}[A_1 \cap A_2] = v_1\mathcal{H}[A_1] \cap v_2\mathcal{H}[A_1].$$

Since (D, u) is some support of \mathcal{B}_1 , we have $D_1 \subseteq D$ and $u_1\mathcal{H}[D_1] \subseteq u\mathcal{H}[D]$. Hence there is a D -path $u_1 \rightarrow u$ labelled k , say, which runs inside A_1 , and a D -path $u \rightarrow s$ labelled m . Altogether, there is a D -path $s_1 \rightarrow s$ labelled m_1km (this path runs entirely in $v_1\mathcal{H}[A_1]$). Since $s_1, s \in \mathcal{B}_1$, there is also a B_1 -path $s_1 \rightarrow s$ where $B_1 = B \cap A_1$, labelled p , say. Again, this path runs inside $v_1\mathcal{H}[A_1]$. Since $H[A_1]$ is retractable, we have $[p]_{H[A_1]} = [p']_{H[A_1]}$ where p' is the word obtained from p by deletion of all letters not in D . Hence there is a D -path $s_1 \rightarrow s$ which runs entirely in $\mathcal{B}_1 \cap u\mathcal{H}[D]$ and, in particular, $s_1 \in u\mathcal{H}[D] \subseteq u\mathcal{H}[A_1 \cap A_2] = v_1\mathcal{H}[A_1] \cap v_2\mathcal{H}[A_2]$ so that $s_1 \in \mathcal{B}_1 \cap \mathcal{B}_2$. Since (D_1, u_1) supports s_1 and therefore also $\mathcal{B}_1 \cap \mathcal{B}_2$, it follows $D \subseteq D_1$ and therefore $D = D_1$ as the converse inclusion has been already shown. In particular, (D, u) provides unique minimal support of \mathcal{B}_1 which is attained at $s_1 \in \mathcal{B}_1 \cap \mathcal{B}_2$. So the configuration in Figure 7 really looks as depicted in Figure 8. By the same reasoning we obtain that $s_2 \in \mathcal{B}_1 \cap \mathcal{B}_2$ and $D_2 = D$.




FIGURE 8

Altogether, $s_1, s_2 \in \mathcal{B}_1 \cap \mathcal{B}_2$ and (D, u) provides unique minimal support of \mathcal{B}_1 as well as \mathcal{B}_2 , attained at s_1 as well as s_2 . Now we continue by induction. Let $2 \leq k < n$ and suppose, subject to some renumbering of the cosets \mathcal{B}_i , we have already shown that $s_1, \dots, s_k \in \mathcal{B}_1 \cap \dots \cap \mathcal{B}_k$ and all these \mathcal{B}_i have unique minimal support (D, u) attained at all these s_i . Again there are $j \in \{1, \dots, k\}$ and $i \in \{k+1, \dots, n\}$ such that $\mathcal{B}_j \cap \mathcal{B}_i \neq \emptyset$ and after some renumbering we may assume that $j = k$ and $i = k+1$. Then, as for the case $k = 1$, $s_k, s_{k+1} \in \mathcal{B}_k \cap \mathcal{B}_{k+1}$ and the unique minimal support of $\mathcal{B}_k \cap \mathcal{B}_{k+1}$ is (D, u) . Again, $s_k \in \mathcal{B}_1 \cap \dots \cap \mathcal{B}_k \cap \mathcal{B}_{k+1}$ and so $\mathcal{B}_j \cap \mathcal{B}_{k+1} \neq \emptyset$ for all $j \leq k$, therefore $s_j, s_{k+1} \in \mathcal{B}_j \cap \mathcal{B}_{k+1}$ and hence $s_1, \dots, s_{k+1} \in \mathcal{B}_1 \cap \dots \cap \mathcal{B}_{k+1}$ and (D, u) provides unique minimal support for \mathcal{B}_{k+1} attained at s_{k+1} . So $s_1, \dots, s_n \in \mathcal{B}_1 \cap \dots \cap \mathcal{B}_n$ and $\bigcup_{i=1}^n \mathcal{B}_i$ has unique minimal support (D, u) attained at some vertices of $\bigcap_{i=1}^n \mathcal{B}_i$.

It remains to argue that \mathcal{B} is indeed a B -cluster. From $\bigcap_{i=1}^n \mathcal{B}_i \neq \emptyset$ we have in particular that $\bigcap_{i=1}^n v_i\mathcal{H}[A_i] \neq \emptyset$. By induction and using (3.10) and Lemma 3.21 we can show that $\bigcap_{i=1}^n v_i\mathcal{H}[A_i] = w\mathcal{H}[C]$ for some vertex $w \in \mathcal{L}$ and $C = \bigcap_{i=1}^n A_i$. From the definition of $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$ it follows that

the graph $\bigcup_{i=1}^n v_i \mathcal{H}[A_i] = \bigcup_{i=1}^n w \mathcal{H}[A_i]$ is isomorphic with the A -cluster $\text{CL}(H[A], \{A_1, \dots, A_n\})$. Corollary 3.12 now implies that $\mathcal{B} = \bigcup_{i=1}^n \mathcal{B}_i$ is isomorphic with the B -cluster $\text{CL}(H[B], \{B \cap A_1, \dots, B \cap A_n\})$. \square

The case $|A| = 2$, which is not handled in Theorem 4.2, is actually trivial.

Proposition 4.3. *Let H be an E -group, $A \subseteq E$ with $|A| = 2$ and $H[A]$ be retractable. Then every connected A -subgraph \mathcal{L} of $\mathcal{H}[A]$ is admissible for $\preceq A$ -coset extension and the full $\preceq A$ -coset extension $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$ has the cluster property.*

Proof. Definition 3.22 is fulfilled for trivial reasons: only the empty set $C = \emptyset$ satisfies $C \subsetneq B \subsetneq A$. Every constituent coset of $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$ is of the form $v \mathcal{H}[a]$ for some letter $a \in A$. Hence, for $B \subsetneq A$, the only B -components of $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$ which have empty intersection with \mathcal{L} are singleton vertices which clearly have unique minimal support. \square

Combination of this with Theorem 4.2 implies the next result.

Corollary 4.4. *Let H be an E -group, $A \subseteq E$, $|A| \geq 3$ and suppose that $H[A]$ is retractable. Let $\mathcal{L} \subseteq \mathcal{H}[A]$ be a connected A -graph which is admissible for $\preceq A$ -coset extension. Assume that for all $B \subsetneq A$ and every vertex $v \in \mathcal{L}$ the full $\preceq B$ -coset extension $\text{CE}(H, v \mathcal{L}[B]; \mathbb{P}_B)$ embeds into $v \mathcal{H}[B]$ and is bridge-free; then the full $\preceq A$ -coset extension $\text{CE}(H, \mathcal{L}; \mathbb{P}_A)$ has the cluster property.*

Proof. In order to reduce the claim of the corollary to Theorem 4.2, we merely need to argue that the graphs $\text{CE}(H, v \mathcal{L}[B]; \mathbb{P}_B)$ for $B \subsetneq A$ have the cluster property. This is proved by induction on $|A|$. For $|A| = 3$ we only need to consider $|B| = 2$, so that $\text{CE}(H, v \mathcal{L}[B]; \mathbb{P}_B)$ has the cluster property by Proposition 4.3. For $|A| > 3$ we can use the inductive claim for all $|B| < |A|$ (in the rôle of A) to find that $\text{CE}(H, v \mathcal{L}[B]; \mathbb{P}_B)$ has the cluster property. \square

For the following recall Definition 3.1 of when a Cayley graph \mathcal{G} covers a graph \mathcal{C} (in terms of canonical morphisms), Definition 3.2 of a k -retractable group and Definition 3.7 of a k -stable expansion.

Definition 4.5. Suppose that a Cayley graph \mathcal{G} covers a complete connected graph \mathcal{C} via a canonical morphism $\varphi: \mathcal{G} \twoheadrightarrow \mathcal{C}$ and let $\mathcal{L} \subseteq \mathcal{C}$ be a connected subgraph. A *cover of \mathcal{L} in \mathcal{G}* (a \mathcal{G} -cover for short) is any connected component of the graph $\varphi^{-1}(\mathcal{L}) \subseteq \mathcal{G}$.

Recall that a crucial feature of covers is the *path lifting property*: if \mathcal{L} admits a path $u \rightarrow v$ labelled $p \in \tilde{E}^*$ and u' is any vertex of $\varphi^{-1}(\mathcal{L})$ such that $\varphi(u') = u$, then $\varphi^{-1}(\mathcal{L})$ admits a path labelled p with initial vertex u' that maps onto the original path in \mathcal{L} under φ .

Theorem 4.6 (upward induction). *Let $1 \leq k < |E|$ and let H be an E -group which is $(k + 1)$ -retractable. Let $A \subseteq E$ with $|A| = k + 1$ and let \mathcal{L}_H be a connected A -subgraph of $\mathcal{H}[A]$ such that*

- (1) \mathcal{L}_H is admissible for $\vartriangleleft A$ -coset extension (with respect to H),
- (2) the full $\vartriangleleft A$ -coset extension $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ has the cluster property.

Let $G \twoheadrightarrow H$ be a k -stable expansion of E -groups such that the Cayley graph \mathcal{G} of G covers all graphs of the form $\overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A) \circledast \mathcal{H}[B]}$ for $B \subsetneq A$ and v a vertex of $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ (thus, in particular, \mathcal{G} covers the graph $\overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ itself). Let \mathcal{L}_G be any cover of \mathcal{L}_H in \mathcal{G} . Then the following hold:

- (i) \mathcal{L}_G is admissible for $\vartriangleleft A$ -coset extension (with respect to G),
- (ii) the full $\vartriangleleft A$ -coset extension $\text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ embeds into $\mathcal{G}[A]$,
- (iii) the full $\vartriangleleft A$ -coset extension $\text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ is bridge-free in $\mathcal{G}[A]$.

Proof. As for (i), that \mathcal{L}_G is admissible for $\vartriangleleft A$ -coset extension follows from the fact that \mathcal{L}_H is admissible for $\vartriangleleft A$ -coset extension and that the canonical morphism $G \twoheadrightarrow H$ is k -stable. In this case, the canonical morphism $\varphi: \mathcal{G} \twoheadrightarrow \mathcal{H}$ is injective on B -components for $B \subsetneq A$ so that condition (3.5) is satisfied for \mathcal{L}_G if it is satisfied for $\mathcal{L}_H = \varphi(\mathcal{L}_G)$.

Towards injectivity as required for (ii), let $\psi: \text{CE}(G, \mathcal{L}_G; \mathbb{P}_A) \rightarrow \mathcal{G}[A]$ be the canonical graph morphism of Proposition 3.18. We first show that for every $B \subsetneq A$ the restriction of ψ to $\text{CE}(G, \mathcal{L}_G; B)$ is injective. Suppose this were not the case. Since the restriction to \mathcal{L}_G is an embedding, that could only happen if two vertices of two distinct constituent cosets $u\mathcal{G}[B]$ and $v\mathcal{G}[B]$ of $\text{CE}(G, \mathcal{L}_G; B)$ were mapped to the same vertex of $\mathcal{G}[A]$ and therefore the cosets $u\mathcal{G}[B]$ and $v\mathcal{G}[B]$ coincide as cosets of $\mathcal{G}[A]$ (see the discussion leading to Remark 3.19). The result in $\mathcal{G}[A]$ is depicted in Figure 9 (left-hand side). By assumption, \mathcal{G} covers the graph $\overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$; so there is




FIGURE 9

a canonical graph morphism $\varphi: \mathcal{G} \twoheadrightarrow \overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ mapping \mathcal{L}_G onto \mathcal{L}_H . Since the expansion $G \twoheadrightarrow H$ is k -stable and $|B| \leq k$, the morphism φ maps $u\mathcal{G}[B] = v\mathcal{G}[B]$ isomorphically onto $\varphi(u)\mathcal{H}[B]$ and likewise onto $\varphi(v)\mathcal{H}[B]$. Hence $\varphi(u)\mathcal{H}[B] = \varphi(v)\mathcal{H}[B]$ in $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$, so that $\varphi(u)$ and $\varphi(v)$ are in the same B -component of \mathcal{L}_H . It follows that $\varphi(u)$ and $\varphi(v)$ can be connected by a B -path which runs in \mathcal{L}_H . Under φ that path lifts to a path $u \rightarrow v'$ which runs in $\mathcal{L}_G \cap u\mathcal{G}[B]$. In particular, $v' \in u\mathcal{G}[B] = v\mathcal{G}[B]$ and $\varphi(v') = \varphi(v)$. Since φ is injective on B -cosets, $v' = v$ and therefore u and v belong to the same B -component of \mathcal{L}_G . It follows that the constituent cosets $u\mathcal{G}[B]$ and $v\mathcal{G}[B]$ of $\text{CE}(G, \mathcal{L}_G; B)$ coincide.

So, for the injectivity claim of (ii), it remains to consider the case when vertices of distinct coset extension $\text{CE}(G, \mathcal{L}_G; B)$ and $\text{CE}(G, \mathcal{L}_G; C)$ would violate injectivity. Let $B, C \subsetneq A$, $B \neq C$ and $x \in \text{CE}(G, \mathcal{L}_G; B)$, $y \in \text{CE}(G, \mathcal{L}_G; C)$ be vertices such that $\psi(x) = \psi(y)$. We need to show that $x = y$ in $\text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ (that is, x and y both are in $\text{CE}(G, \mathcal{L}_G; B \cap C)$ and coincide). From $\psi(x) = \psi(y)$ we see that in $\mathcal{G}[A]$ the situation is as depicted in Figure 9 (right-hand side) with $\psi(x) = z = \psi(y)$. That is, u and z are connected by a B -path while v and z are connected by a C -path, so that $z \in u\mathcal{G}[B] \cap v\mathcal{G}[C]$. Let us consider some canonical graph morphism $\mathcal{G} \rightarrow \overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ (according to the statement of the Theorem), which maps \mathcal{L}_G onto \mathcal{L}_H . Let u', v', z' be the image vertices of u, v, z , respectively, under this morphism. Then $u', v' \in \mathcal{L}_H$ and $z' \in u'\mathcal{H}[B] \cap v'\mathcal{H}[C]$. The latter intersection is a $(B \cap C)$ -(constituent) coset of $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$, having non-empty intersection with the skeleton \mathcal{L}_H , say $u'\mathcal{H}[B] \cap v'\mathcal{H}[C] = c\mathcal{H}[B \cap C]$ for some $c \in \mathcal{L}_H$. Moreover, the intersections $\mathcal{L}_H \cap u'\mathcal{H}[B]$ and $\mathcal{L}_H \cap v'\mathcal{H}[C]$ both are connected (namely B - respectively C -components of \mathcal{L}_H). The situation is depicted in Figure 10. So there are paths $u' \xrightarrow{p} c$ in $\mathcal{L}_H \cap u'\mathcal{H}[B]$,




FIGURE 10

$v' \xrightarrow{q} c$ in $\mathcal{L}_H \cap v'\mathcal{H}[C]$ and $c \xrightarrow{r} z'$ in $u'\mathcal{H}[B] \cap v'\mathcal{H}[C]$. In particular, pr labels a path $u' \rightarrow z'$, qr labels a path $v' \rightarrow z'$. From k -stability of the expansion $G \rightarrow H$ it follows that the morphism $\mathcal{G} \rightarrow \overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ is injective on all cosets $x\mathcal{G}[D]$ for all $D \subsetneq A$. In particular, this morphism is bijective between $u\mathcal{G}[B]$ and $u'\mathcal{H}[B]$ as well as between $v\mathcal{G}[C]$ and $v'\mathcal{H}[C]$. From this it follows that the paths in $u'\mathcal{H}[B] \cup v'\mathcal{H}[C]$ just mentioned lift to paths in $u\mathcal{G}[B] \cup v\mathcal{G}[C]$: hence there is a path $u \rightarrow z$ labelled pr and one $v \rightarrow z$ labelled qr . It follows that, in $\text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$,

$$u \cdot p = z \cdot r^{-1} = v \cdot q.$$

Since $p: u' \rightarrow c$ runs in \mathcal{L}_H and so does $q: v' \rightarrow c$, the path $p: u \rightarrow z \cdot r^{-1}$ runs in \mathcal{L}_G , and so does the path $q: v \rightarrow z \cdot r^{-1}$. It follows that

$$u\mathcal{G}[B] = (z \cdot r^{-1})\mathcal{G}[B] \text{ and } v\mathcal{G}[C] = (z \cdot r^{-1})\mathcal{G}[C],$$

thus $u\mathcal{G}[B] \cap v\mathcal{G}[C] = (z \cdot r^{-1})\mathcal{G}[B \cap C]$ so that, in $\text{CE}(G, \mathcal{L}_G; \{B, C\})$:

$$x = (z \cdot r^{-1}) \cdot r = y,$$

that is, x and y represent the same vertex in $\text{CE}(G, \mathcal{L}_G; B \cap C)$, as required. Altogether, $\text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ embeds in $\mathcal{G}[A]$ via the morphism of Proposition 3.18.

It remains to argue for (iii), that $\text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ is bridge-free. So we look at a pair of vertices $v_1, v_2 \in \mathcal{L}_G$, subsets $A_1, A_2 \subsetneq A$, and vertices $s_1 \in v_1\mathcal{G}[A_1]$, $s_2 \in v_2\mathcal{G}[A_2]$, and assume that, for some $B \subsetneq A$, there is a B -path $s_1 \xrightarrow{p} s_2$ running in $\mathcal{G}[A]$ (all the following takes place in $\mathcal{G}[A]$ as depicted in Figure 11). In addition, there are an A -path $v_1 \xrightarrow{q} v_2$ running in




FIGURE 11

\mathcal{L}_G and A_i -paths $v_i \xrightarrow{f_i} s_i$ running in $v_i\mathcal{G}[A_i]$. Consider the canonical graph morphism $\varphi: \mathcal{G} \rightarrow \overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$, which maps \mathcal{L}_G onto \mathcal{L}_H . Let v'_i be the image of v_i in \mathcal{L}_H under this morphism. The path $v_1 \xrightarrow{q} v_2$ is mapped to the path $v'_1 \xrightarrow{q} v'_2$ in \mathcal{L}_H . Let us denote the image of s_i by s'_i ; then the path $v_i \xrightarrow{f_i} s_i$ running in $v_i\mathcal{G}[A_i]$ is mapped to the path $v'_i \xrightarrow{f_i} s'_i$ which runs in $v'_i\mathcal{H}[A_i]$. So far, all these paths run in $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$. Further, the path $s_1 \xrightarrow{p} s_2$ is mapped to the path $s'_1 \xrightarrow{p} s'_2$, which runs in $\overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$. It follows that there is a B -path $s'_1 \xrightarrow{p^\circ} s'_2$ which runs in $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ (in fact, p° is the word obtained from p by deletion of the letters which traverse loop edges of $\overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)} \setminus \text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$).

So consider the B -component \mathcal{B} of $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ which contains the two vertices s'_1 and s'_2 . The cluster property of $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ shows the following: either \mathcal{B} has non-empty intersection with the skeleton \mathcal{L}_H , or else \mathcal{B} is a B -cluster (the existence of unique minimal support is not needed in this context). Assume the latter case first: as a B -cluster, \mathcal{B} is the union $\mathcal{B} = \mathcal{B}_1 \cup \dots \cup \mathcal{B}_n$ of $(B \cap C_i)$ -cosets where $C_i \subsetneq A$, $|C_i| = |A| - 1$ and assume first that $n \geq 2$; the case $n = 1$ will be handled below. We may assume that $s'_i \in \mathcal{B}_i$ for $i = 1, 2$. The pairs (A_1, v'_1) and (A_2, v'_2) provide support for s'_1 and s'_2 , respectively. The cosets $\mathcal{B}_1 = s'_1\mathcal{H}[C_1 \cap B] \subseteq v'_1\mathcal{H}[C_1]$ and $\mathcal{B}_2 = s'_2\mathcal{H}[C_2 \cap B] \subseteq v'_2\mathcal{H}[C_2]$ have non-empty intersection (indeed,

$\mathcal{B}_1 \cap \mathcal{B}_2$ contains the core of \mathcal{B}). Hence $v'_1\mathcal{H}[C_1] \cap v'_2\mathcal{H}[C_2] \neq \emptyset$ so that $v'_1\mathcal{H}[C_1] \cap v'_2\mathcal{H}[C_2] = v\mathcal{H}[C]$ for $C = C_1 \cap C_2$ and some vertex $v \in \mathcal{L}_H$. The situation is depicted in Figure 12. In particular, there is a vertex




FIGURE 12

$s \in s'_1\mathcal{H}[B \cap C_1] \cap s'_2\mathcal{H}[B \cap C_2]$ and there are $B \cap C_i$ -paths

$$s'_1 \xrightarrow{p_1} s \xrightarrow{p_2} s'_2$$

labelled p_i ($i = 1, 2$). We now consider the B -augmentation of $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ at the vertex s and the canonical graph morphism

$$\psi: \mathcal{G} \rightarrow \overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A) \otimes \mathcal{H}[B]}$$

which maps the covering graph \mathcal{L}_G onto \mathcal{L}_H . The graphs $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ and $\overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A) \otimes \mathcal{H}[B]}$ are almost the same except that the cluster \mathcal{B} in the coset extension $\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ is blown up to the full coset $s\mathcal{H}[B]$ in the latter graph. The morphism ψ now maps the path $s_1 \xrightarrow{p} s_2$ to the path $s'_1 \xrightarrow{p} s'_2$ which runs in $s\mathcal{H}[B]$; but $s'_1 \xrightarrow{p_1} s \xrightarrow{p_2} s'_2$ also run in $s\mathcal{H}[B]$ which implies that $[p]_H = [p_1 p_2]_H$. Since the expansion $G \twoheadrightarrow H$ is k -stable and $|B| \leq k$, it follows that $[p]_G = [p_1 p_2]_G$. In addition, k -stability implies that ψ provides isomorphisms $v_1\mathcal{G}[C_1] \twoheadrightarrow v'_1\mathcal{H}[C_1]$ and $v_2\mathcal{G}[C_2] \twoheadrightarrow v'_2\mathcal{H}[C_2]$ and therefore also an isomorphism $v_1\mathcal{G}[C_1] \cup v_2\mathcal{G}[C_2] \twoheadrightarrow v'_1\mathcal{H}[C_1] \cup v'_2\mathcal{H}[C_2]$. It follows that the path $s_1 \xrightarrow{p_1} s \cdot p_1$ runs in $v_1\mathcal{G}[C_1]$ while $s_1 \cdot p_1 \xrightarrow{p_2} s_1 \cdot p_1 p_2 = s_2$ runs in $v_2\mathcal{G}[C_2]$. So the path $s_1 \xrightarrow{p_1 p_2} s_2$ runs entirely in $v_1\mathcal{G}[C_1] \cup v_2\mathcal{G}[C_2] \subseteq \text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ and thus provides a B -path between s_1 and s_2 in the coset extension $\text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$. Finally, for the same reason, we see that in case $n = 1$, that is, $\mathcal{B} = \mathcal{B}_1 \subseteq v'_1\mathcal{H}[C_1]$ the path $s_1 \xrightarrow{p_1 p_2} s_2$ runs in $v_1\mathcal{G}[C_1]$ which is contained in $\text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$.

The remaining case, where \mathcal{B} has non-empty intersection with the skeleton \mathcal{L}_H , is easy: in this case \mathcal{B} is a full B -coset $\mathcal{B} = v\mathcal{H}[B]$ for some vertex $v \in \mathcal{L}_H$. The canonical morphism $\varphi: \mathcal{G} \rightarrow \overline{\text{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ induces an isomorphism $\phi: s_1\mathcal{G}[B] = s_2\mathcal{G}[B] \twoheadrightarrow v\mathcal{H}[B]$ where $\phi = \varphi \upharpoonright s_1\mathcal{G}[B]$ is the restriction.

Then $s_1\mathcal{G}[B] = s_2\mathcal{G}[B] = \phi^{-1}(v\mathcal{H}[B]) = \phi^{-1}(v)\mathcal{G}[B]$. But $\phi^{-1}(v) \in \mathcal{L}_G$ so that $s_1\mathcal{G}[B] = s_2\mathcal{G}[B]$ is contained in $\text{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$. \square

5. CONSTRUCTION OF THE GROUP G

The group G announced in Lemma 2.9 will be constructed via a series of expansions

$$G_1 \leftarrow H_1 \leftarrow G_2 \leftarrow \cdots \leftarrow G_{|E|-1} \leftarrow H_{|E|-1} \leftarrow G_{|E|} = G \quad (5.1)$$

where, for every k , the expansions $G_k \leftarrow H_k$ and $H_k \leftarrow G_{k+1}$ are k -stable and the groups H_k and G_{k+1} are $(k+1)$ -retractable. Here the series (3.4) is interleaved with the intermediate stages H_k in (5.1). The series (5.1) is defined by an ascending series

$$\mathcal{X}_1 \subseteq \mathcal{Y}_1 \subseteq \mathcal{X}_2 \subseteq \cdots \subseteq \mathcal{X}_{|E|-1} \subseteq \mathcal{Y}_{|E|-1} \subseteq \mathcal{X}_{|E|} \quad (5.2)$$

of complete E -graphs such that each group in the series (5.1) is the transition group of the corresponding graph in (5.2), that is,

$$G_k = \mathcal{T}(\mathcal{X}_k) \text{ and } H_k = \mathcal{T}(\mathcal{Y}_k)$$

for all k in question. Every graph in the series (5.2) is obtained from its predecessor by adding certain complete components. These components are constructed by an inductive procedure, the idea of which is as follows. The graph \mathcal{X}_1 is obtained as a suitable completion of the given oriented graph $\mathcal{E} = (V \cup \tilde{E}; \alpha, \omega, {}^{-1})$, here considered as an E -labelled graph where every edge gets its own label. This serves to initialise the series (5.1) with $G_1 := \mathcal{T}(\mathcal{X}_1)$.

Suppose that for $k \geq 1$ the graph \mathcal{X}_k and therefore its transition group G_k have already been constructed. Then the step $\mathcal{X}_k \rightsquigarrow \mathcal{Y}_k$, and hence the step $G_k \rightsquigarrow H_k$, raises the ‘‘degree of retractability’’ from k to $k+1$ and thereby lays the ground for the transition $H_k \rightsquigarrow G_{k+1}$. That step is intended to ensure the following: suppose that p is a word over $k+1$ letters which forms a path $u \rightarrow v$ in \mathcal{E} and $a \in \text{co}(p)$ for some $a \in E$; if H_k satisfies the relation $p = p_{a \rightarrow 1}$, but there is no word q in the letters $B := \text{co}(p) \setminus \{a\}$ (and their inverses) forming a path $u \rightarrow v$ in \mathcal{E} such that H_k satisfies the relation $p = q$, then some component of $\mathcal{X}_{k+1} \setminus \mathcal{Y}_k$ guarantees that G_{k+1} avoids the relation $p = p_{a \rightarrow 1}$ and therefore every relation $p = q$ with $q \in \tilde{B}^*$.

5.1. Definition of G_1 and the transition $G_k \rightsquigarrow H_k$. The idea of the construction of the graph \mathcal{X}_1 is to extend the given oriented graph $\mathcal{E} = (V \cup \tilde{E}, \alpha, \omega, {}^{-1})$ to a complete E -graph on the vertex set V in whose transition group the permutation $[e]$ corresponding to any non-loop edge e is the transposition of V that swaps the two vertices ae and ωe . Let $\mathcal{E} = (V \cup \tilde{E}; \alpha, \omega, {}^{-1})$ be a finite connected oriented graph. We let the set of positive edges E be our alphabet and label every edge e by itself. Thereby we get the E -labelled graph $(V \cup \tilde{E}; \alpha, \omega, {}^{-1}, \ell, E)$ where ℓ is the identity function mapping every $e \in \tilde{E}$, considered as an edge, to itself, considered

as a label. The resulting graph is an E -graph for trivial reasons, since every label appears exactly once.

Next, for every non-loop edge e we add a new edge \bar{e} and set

$$\alpha\bar{e} := \omega e, \quad \omega\bar{e} := \alpha e, \quad \ell(\bar{e}) := \ell(e) = e.$$

We have thus completed every non-loop edge $u \xrightarrow{e} v$ to a 2-cycle $u \begin{array}{c} \xleftarrow{e} \\ \xrightarrow{e} \end{array} v$.

Let us denote the set of all positive edges so obtained (the original ones and the added ones) by F ; then the oriented E -graph $\mathcal{F} = (V \cup \bar{F}; \alpha, \omega, {}^{-1}, \ell, E)$ is weakly complete. Let $\mathcal{X}_1 := \bar{\mathcal{F}}$ be its trivial completion. The transition group $G_1 := \mathcal{T}(\mathcal{X}_1)$ is an E -group of permutations acting on the vertex set V . For every $e \in E$, $[e]_{G_1}$ is either a transposition (if e is not a loop edge then $[e]$ swaps αe and ωe) or the identity permutation (if e is a loop edge). Note that two distinct labels $e, f \in E$ may represent the same permutation of V (since we allow multiple edges in \mathcal{E}).

Remark 5.1. Instead of completing all non-loop edges to 2-cycles we could equally well complete every such edge e to an n -cycle for any fixed $n \geq 2$, by attaching to the edge $u \xrightarrow{e} v$ an e -path $u \xleftarrow{e} \cdots \xleftarrow{e} v$ consisting of a sequence of $n - 1$ new edges labelled e and $n - 2$ new intermediate vertices. In the resulting transition group, the permutation $[e]$ assigned to e then is a cyclic permutation of length n mapping αe to ωe . Distinct labels coming from non-loop edges then automatically represent different permutations provided that $n \geq 3$.

The transition from G_k to H_k is easily described. Suppose we have already defined the graph \mathcal{X}_k and thus the group $G_k = \mathcal{T}(\mathcal{X}_k)$. We set

$$\mathcal{Y}_k := \mathcal{X}_k \sqcup \bigsqcup \{ \overline{\mathcal{G}_k[A]} : A \subseteq E, |A| = k \}.$$

Provided that G_k is k -retractable, the transition group $H_k = \mathcal{T}(\mathcal{Y}_k)$ is $(k + 1)$ -retractable and the expansion $G_k \leftarrow H_k$ is k -stable (Theorem 3.8). In particular, H_1 is 2-retractable.

5.2. The transition $H_k \rightsquigarrow G_{k+1}$. The expansion $H_k \leftarrow G_{k+1}$ is more delicate. We assemble a complete E -graph $\mathcal{X}_{k+1} = \mathcal{Y}_k \sqcup \overline{\mathcal{Z}_k}$ to obtain G_{k+1} as the transition group $G_{k+1} = \mathcal{T}(\mathcal{X}_{k+1})$. The new, weakly complete components of \mathcal{Z}_k will be constructed as augmentations of clusters and coset extensions based on H_k . To this end we first collect, for $k \geq 2$, properties of the precursors G_k and H_{k-1} of H_k , which then serve as conditions to be maintained inductively also in the passage to G_{k+1} . At level k , we denote these inductive conditions as COND_k for the pair (G_k, H_{k-1}) . So COND_k will serve as inductive hypothesis for the construction of \mathcal{Z}_k , and hence G_{k+1} , which then needs to guarantee that the conditions COND_{k+1} are satisfied by the pair (G_{k+1}, H_k) . In the following, we identify subgraphs of \mathcal{E} with their labelled versions inside \mathcal{X}_1 .

Condition 5.2. As conditions COND_k , for $k \geq 2$, we collect the following:

- (i) H_{k-1} and G_k are k -retractable and the expansion $G_k \twoheadrightarrow H_{k-1}$ is $(k-1)$ -stable,

and, for every $B \subseteq E$ with $|B| \leq k$, for the \mathcal{G}_k -cover $\mathcal{C}_{\mathcal{G}_k}$ of every connected component \mathcal{C} of $\mathcal{B} = \langle B \rangle$ in $\mathcal{E} \subseteq \mathcal{X}_1$, the following hold:

- (ii) $\mathcal{C}_{\mathcal{G}_k}$ is admissible for \mathcal{B} -coset extension,
- (iii) the full \mathcal{B} -coset extension $\text{CE}(G_k, \mathcal{C}_{\mathcal{G}_k}; \mathbb{P}_B)$ embeds into $\mathcal{G}_k[B]$,
- (iv) the full \mathcal{B} -coset extension $\text{CE}(G_k, \mathcal{C}_{\mathcal{G}_k}; \mathbb{P}_B)$ is bridge-free.

By Theorem 3.8, (i) implies that $H_k \twoheadrightarrow G_k$ in particular is k -stable, as already mentioned in context with the definition of H_k . Let $\psi_k: \mathcal{G}_k \rightarrow \mathcal{X}_1$ be some canonical graph morphism, $\chi_k: \mathcal{H}_k \rightarrow \mathcal{G}_k$ the graph morphism induced by the canonical morphism $H_k \twoheadrightarrow G_k$, and let $\varphi_k = \psi_k \circ \chi_k$. By k -stability, χ_k is injective on connected B -subgraphs for $|B| \leq k$. Let $A \subseteq E$ be a set of $|A| = k+1$ many (positive) edges of $\mathcal{E} \subseteq \mathcal{X}_1$, $\mathcal{A} = \langle A \rangle$ be the subgraph of \mathcal{E} spanned by A . Let \mathcal{C} be a connected component of \mathcal{A} and $\mathcal{C}_{\mathcal{H}_k}$ be an \mathcal{H}_k -cover of \mathcal{C} , that is, some connected component of $\varphi_k^{-1}(\mathcal{C})$. Next we show that $\mathcal{C}_{\mathcal{H}_k}$ is admissible for \mathcal{A} -coset extension (with respect to H_k) and that the full \mathcal{A} -coset extension $\text{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A)$ has the cluster property. From this it will follow that augmented coset extensions of the form $\text{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A) \circledast \mathcal{H}_k[B]$ are well defined; they will be essential ingredients of the graph \mathcal{Z}_k to be defined below (Definition 5.3).

Let $B \subsetneq A$ and let $\mathcal{U} \subseteq \mathcal{C}_{\mathcal{H}_k}$ be some B -component of $\mathcal{C}_{\mathcal{H}_k}$. Then $\mathcal{U}' := \varphi_k(\mathcal{U}) \subseteq \mathcal{C}$ is a B -component of \mathcal{C} and hence is a connected component of $\mathcal{B} := \langle B \rangle$. By the inductive hypothesis, any \mathcal{G}_k -cover $\mathcal{U}'_{\mathcal{G}_k}$ of \mathcal{U}' is admissible for \mathcal{B} -coset extension (with respect to G_k) and $\text{CE}(G_k, \mathcal{U}'_{\mathcal{G}_k}; \mathbb{P}_B)$ embeds into $\mathcal{G}_k[B]$ and is bridge-free. Since the morphism $\chi_k: \mathcal{H}_k \rightarrow \mathcal{G}_k$ is injective on B -components (that is, injective on B -cosets), it follows that $\mathcal{U}'_{\mathcal{G}_k} \cong \mathcal{U}$ and hence also

$$\text{CE}(G_k, \mathcal{U}'_{\mathcal{G}_k}; \mathbb{P}_B) \cong \text{CE}(H_k, \mathcal{U}; \mathbb{P}_B). \quad (5.3)$$

As the latter graph embeds into $\mathcal{H}_k[B] \cong \mathcal{G}_k[B]$, it follows that condition (3.5) of Definition 3.16 is fulfilled so that $\mathcal{C}_{\mathcal{H}_k}$ is admissible for \mathcal{A} -coset extension (with respect to H_k). Once more by the inductive hypothesis, every graph in (5.3) is bridge-free. Then, by Corollary 4.4, the full \mathcal{A} -coset extension $\text{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A)$ itself has the cluster property.

We can now define the components of the graph \mathcal{Z}_k .

Definition 5.3. The graph \mathcal{Z}_k is the disjoint union of

- (1) all augmented A -clusters

$$\text{CL}(H_k[A], \mathbb{P}) \circledast \mathcal{H}_k[B]$$

for $A \subseteq E$ with $|A| = k+1$, \mathbb{P} a set of proper subsets of A , v a vertex of $\text{CL}(H_k[A], \mathbb{P})$ and $B \subsetneq A$;

- (2) all augmented full \mathcal{A} -coset extensions

$$\text{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A) \circledast \mathcal{H}_k[B]$$

for $A \subseteq E$ with $|A| = k + 1$, \mathcal{C} a connected component of $\mathcal{A} = \langle A \rangle$, $\mathcal{C}_{\mathcal{H}_k}$ an \mathcal{H}_k -cover of \mathcal{C} , \mathbb{P}_A the set of all proper subsets of A , v a vertex of $\mathbf{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A)$ and $B \subsetneq A$.

We note that the augmented clusters and augmented coset extensions contain, for $B = \emptyset$, all “plain” clusters and coset extensions. Recall that $G_{k+1} = \mathcal{T}(X_{k+1}) = \mathcal{T}(Y_k \sqcup \overline{Z}_k)$.

Proposition 5.4. *The expansion $H_k \leftarrow G_{k+1}$ is k -stable and hence G_{k+1} is $(k + 1)$ -retractable.*

Proof. We need to prove k -stability, the second assertion then follows from Theorem 3.8 by inductive hypothesis (i) and the definition of H_k . Let $C \subseteq E$ with $|C| = k$, let $p \in \tilde{C}^*$ and assume that $[p]_{G_{k+1}} \neq 1$; we need to show that $[p]_{H_k} \neq 1$. There exists a component \mathcal{L} of Y_k or of \overline{Z}_k witnessing the inequality $[p]_{G_{k+1}} \neq 1$. That is, in this component there is a vertex v such that $v \cdot p \neq v$. If the witnessing component \mathcal{L} belongs to Y_k , then we are done since then $[p]_{H_k} \neq 1$ immediately follows from $H_k = \mathcal{T}(Y_k)$. If \mathcal{L} is a component of \overline{Z}_k , then $\mathcal{L} = \overline{\mathcal{M}}$ where \mathcal{M} is of the form (1) or (2) in Definition 5.3, and the path $p: v \rightarrow v \cdot p$ runs in the C -component $v\overline{\mathcal{M}}[C]$. Recall that $v\overline{\mathcal{M}}[C]$ denotes the C -component of v in the graph $\overline{\mathcal{M}}$ while $\overline{v\mathcal{M}[C]}$ is the trivial completion of $v\mathcal{M}[C]$, that is, the trivial completion of the C -component of v in \mathcal{M} . Obviously

$$v\mathcal{M}[C] \subseteq v\overline{\mathcal{M}}[C] \subseteq \overline{v\mathcal{M}[C]},$$

and the latter two graphs differ only in loop edges having labels not in C . Hence C -paths in $v\overline{\mathcal{M}}[C]$ and $\overline{v\mathcal{M}[C]}$ traverse the same edges and meet the same vertices. It is therefore sufficient to look at $\overline{v\mathcal{M}[C]}$ instead of $v\overline{\mathcal{M}}[C]$. From Corollaries 3.12, 3.15 and Proposition 3.24, and since the (plain) coset extensions in Definition 5.3 (2) have the cluster property, it follows that, for the graph \mathcal{M} in question, the C -component $v\mathcal{M}[C]$ must be isomorphic to one of the following:

- (i) a full C -coset $\mathcal{H}_k[C]$, or
- (ii) a C -cluster $\mathbf{CL}(H_k[C], \mathbb{P})$ for some set \mathbb{P} of proper subsets of C (this includes, for $\mathbb{P} = \{B\}$, also B -cosets $\mathcal{H}_k[B]$ for $B \subsetneq C$), or
- (iii) a D -augmented C -cluster $\mathbf{CL}(H_k[C], \mathbb{P}) @ \mathcal{H}_k[D]$ for some set \mathbb{P} of proper subsets of C , some vertex u of $\mathbf{CL}(H_k[C], \mathbb{P})$ and some proper subset D of C .

In case (i), $v\mathcal{M}[C] \cong \mathcal{H}_k[C]$, so the claim $[p]_{H_k} \neq 1$ again follows immediately. In case (ii) we get

$$v\mathcal{M}[C] \cong \mathbf{CL}(H_k[C], \mathbb{P}) \cong \mathbf{CL}(G_k[C], \mathbb{P}) \cong \mathbf{CL}(H_{k-1}[C], \mathbb{P})$$

where the second isomorphism is obvious since $H_k[C] \cong G_k[C]$ by k -stability of $H_k \rightarrow G_k$ while the third isomorphism follows from Lemma 3.10. In

case (iii) we get

$$\begin{aligned} v\mathcal{M}[C] &\cong \text{CL}(H_k[C], \mathbb{P}) \textcircled{u} \mathcal{H}_k[D] \\ &\cong \text{CL}(G_k[C], \mathbb{P}) \textcircled{t} \mathcal{G}_k[D] \cong \text{CL}(H_{k-1}[C], \mathbb{P}) \textcircled{s} \mathcal{H}_{k-1}[D] \end{aligned}$$

where t and s are the images of u under the canonical morphisms $H_k \rightarrow G_k$ and $H_k \rightarrow H_{k-1}$, respectively, and, again, the second isomorphism is obvious since $H_k[C] \cong G_k[C]$ and $\mathcal{H}_k[D] \cong \mathcal{G}_k[D]$ by k -stability of $H_k \rightarrow G_k$ while the third isomorphism follows from Lemma 3.14. Hence, in cases (ii) and (iii), $v\mathcal{M}[C]$ is isomorphic with a component of $\overline{\mathcal{Z}_{k-1}}$ so that $[p]_{G_k} \neq 1$, from which again $[p]_{H_k} \neq 1$ follows. \square

From Theorem 4.6 it follows that for every set $A \subseteq E$ with $|A| = k + 1$ and every connected component \mathcal{C} of \mathcal{A} , every \mathcal{G}_{k+1} -cover $\mathcal{C}_{\mathcal{G}_{k+1}}$ (that is, every connected component of $\psi_{k+1}^{-1}(\mathcal{C})$ in \mathcal{G}_{k+1} where $\psi_{k+1}: \mathcal{G}_{k+1} \rightarrow \mathcal{X}_1$ is a canonical graph morphism) is admissible for $\mathcal{C}A$ -coset extension, and the full $\mathcal{C}A$ -coset extension $\text{CE}(G_{k+1}, \mathcal{C}_{\mathcal{G}_{k+1}}; \mathbb{P}_A)$ embeds into $\mathcal{G}_{k+1}[A]$ and is bridge-free. If $|A| = l < k + 1$ we have by induction that, for every connected component \mathcal{C} of \mathcal{A} the full $\mathcal{C}A$ -coset extension $\text{CE}(G_l, \mathcal{C}_{\mathcal{G}_l}; \mathbb{P}_A)$ embeds into $\mathcal{G}_l[A]$. But the expansion $G_l \leftarrow G_{k+1}$ is l -stable whence $\text{CE}(G_l, \mathcal{C}_{\mathcal{G}_l}; \mathbb{P}_A) \cong \text{CE}(G_{k+1}, \mathcal{C}_{\mathcal{G}_{k+1}}; \mathbb{P}_A)$ and $\mathcal{G}_l[A] \cong \mathcal{G}_{k+1}[A]$. We have thus maintained Condition 5.2 in the passage from k to $k + 1$ by having verified COND_{k+1} :

- (i) H_k and G_{k+1} are $(k + 1)$ -retractable and the expansion $G_{k+1} \rightarrow H_k$ is k -stable (by Proposition 5.4)

and, for every $A \subseteq E$ with $|A| \leq k + 1$, for the \mathcal{G}_{k+1} -cover $\mathcal{C}_{\mathcal{G}_{k+1}}$ of every connected component \mathcal{C} of $\mathcal{A} = \langle A \rangle$ in \mathcal{E} , the following hold:

- (ii) $\mathcal{C}_{\mathcal{G}_{k+1}}$ is admissible for $\mathcal{C}A$ -coset extension,
- (iii) the full $\mathcal{C}A$ -coset extension $\text{CE}(G_{k+1}, \mathcal{C}_{\mathcal{G}_{k+1}}; \mathbb{P}_A)$ embeds into $\mathcal{G}_{k+1}[A]$,
- (iv) the full $\mathcal{C}A$ -coset extension $\text{CE}(G_{k+1}, \mathcal{C}_{\mathcal{G}_{k+1}}; \mathbb{P}_A)$ is bridge-free.

In addition, the base case for this inductive procedure also works. The group H_1 is 2-retractable and so is G_2 since $G_2 \rightarrow H_1$ is 1-stable. By Proposition 4.3, for every set $A \subseteq E$ with $|A| = 2$, every \mathcal{H}_1 -cover $\mathcal{C}_{\mathcal{H}_1}$ of every component \mathcal{C} of \mathcal{A} is admissible for $\mathcal{C}A$ -coset extension (with respect to H_1) and $\text{CE}(H_1, \mathcal{C}_{\mathcal{H}_1}; \mathbb{P}_A)$ has the cluster property. Theorem 4.6 then implies that the \mathcal{G}_2 -cover $\mathcal{C}_{\mathcal{G}_2}$ is admissible for $\mathcal{C}A$ -coset extension (with respect to G_2) and that $\text{CE}(G_2, \mathcal{C}_{\mathcal{G}_2}; \mathbb{P}_A)$ embeds in $\mathcal{G}_2[A]$ and is bridge-free (the assertions for G_2 can also be checked by direct inspection). In other words, we have shown that conditions COND_2 are satisfied by the pair (G_2, H_1) . Altogether the series of expansions

$$G_1 \leftarrow H_1 \leftarrow G_2 \leftarrow \cdots \leftarrow G_{|E|-1} \leftarrow H_{|E|-1} \leftarrow G_{|E|}$$

is well defined and $G = G_{|E|}$ is retractable.

5.3. Properties of $G = G_{|E|}$. We need to argue that G satisfies the requirements of Lemma 2.9. Requirement (2), that G is retractable, and therefore has a content function by Proposition 3.5, has already been proved.

We are left with showing requirements (1) and (3):

- (1) that every permutation of E induced by an automorphism of \mathcal{E} extends to an automorphism of G , and
- (3) that for every word which forms a path $u \rightarrow v$ in \mathcal{E} there is a G -equivalent word which also forms a path $u \rightarrow v$ and uses only edges of the (common) G -content, or $u = v$ in case of empty content.

We start with item (1); (3) will then be dealt with in Lemma 5.6 and Corollary 5.7. In the context of (1), “an automorphism of \mathcal{E} ” refers to any automorphism of the *unlabelled oriented graph* $\mathcal{E} = (V \cup \tilde{E}; \alpha, \omega, -^1)$. Recall from the definition of an automorphism of an oriented graph that every such automorphism of \mathcal{E} is required to induce a permutation on the set E of positive edges of \mathcal{E} , hence induces a permutation on our labelling alphabet E . Similarly, “an automorphism of G ” means automorphism of the mere group G (rather than of G as an E -group, which cannot have non-trivial automorphisms).

Proposition 5.5. *Every permutation $E \rightarrow E$ induced by an automorphism of the oriented graph \mathcal{E} extends to an automorphism of G .*

Proof. Let γ be a permutation of E induced by an automorphism of \mathcal{E} , also denoted γ . We demonstrate the required property for all G_k and H_k , by induction on k . First note that γ (uniquely) extends to an automorphism $\hat{\gamma}$ of \mathcal{X}_1 from which the claim follows for the group G_1 . Indeed, for every pair of vertices $u, v \in \mathcal{X}_1$ and every word $p \in \tilde{E}^*$, we have $p: u \rightarrow v$ if and only if $\gamma p: \hat{\gamma}u \rightarrow \hat{\gamma}v$. Consequently, for every word $p \in \tilde{E}^*$, G_1 satisfies the relation $p = 1$ if and only if it satisfies $\gamma p = 1$.

So, let $k \geq 1$ and assume inductively that γ extends to an automorphism $\hat{\gamma}$ of \mathcal{X}_k (this means that there is an automorphism $\hat{\gamma}$ of the oriented graph \mathcal{X}_k such that for every edge $e \in \mathcal{X}_k$ we have $\ell(\hat{\gamma}e) = \gamma\ell(e)$); by the same reasoning as for $k = 1$ we see that in this case γ extends to an automorphism of G_k . From the definition of the graph \mathcal{Y}_k it now follows that γ extends to an automorphism $\hat{\gamma}$ of \mathcal{Y}_k which again implies that γ extends to an automorphism of H_k . From this in turn it follows that γ extends to an automorphism of \mathcal{X}_{k+1} and therefore again to an automorphism of G_{k+1} . \square

The assertion of the last proposition is essentially a direct consequence of the fact that the entire process behind our construction of G , on the basis of the given oriented graph \mathcal{E} , is symmetry-preserving. Indeed, none of the intermediate steps involves any choices that could possibly break symmetries in the input data, i.e. could be incompatible with isomorphisms between oriented input graphs \mathcal{E} . In particular, the inductive construction steps reflected in Theorems 4.2 and 4.6 proceed by cardinality of subsets of E and treat all subsets of the same size uniformly and in parallel.² Any

²This should be contrasted e.g. with constructions based on some enumeration of the subsets of E , which could well break symmetries.

isomorphism between oriented graphs $\mathcal{E} \cong \mathcal{E}'$ would successively extend to isomorphisms between the associated graphs $\mathcal{X}_i \cong \mathcal{X}'_i$ and $\mathcal{Y}_i \cong \mathcal{Y}'_i$ and induced isomorphisms between their transition groups $G_i \cong G'_i$ and $H_i \cong H'_i$. In this sense, the entire inductive process underlying the expansion chain (5.1) is isomorphism-respecting, hence in particular compatible with permutations of E stemming from automorphisms of \mathcal{E} .

Finally, we have to deal with requirement (3) of Lemma 2.9. Recall that for a word $p \in \tilde{E}^*$, $\text{co}(p)$ is the set of all letters $a \in E$ for which a or a^{-1} occurs in p . The following lemma is crucial for establishing (3).

Lemma 5.6. *Let $p \in \tilde{E}^*$ be a word that forms a path $u \rightarrow v$ in \mathcal{E} ; let $A = \text{co}(p)$ and suppose that for some letter $a \in A$ and $B = A \setminus \{a\}$ there exists a word $r \in \tilde{B}^*$ such that $[p]_G = [r]_G$. Then there exists a word $q \in \tilde{B}^*$ such that $[p]_G = [q]_G$ and, in addition, q forms a path $u \rightarrow v$ in \mathcal{E} .*

Proof. First recall that every loop edge e of \mathcal{E} induces the identity permutation on the set V of vertices of \mathcal{X}_1 , whence $[e]_{G_1} = 1$; then $[e]_G = 1$ follows from the fact that the expansion $G \twoheadrightarrow G_1$ is 1-stable. Hence, if p contains only loop edges then $u = v$, the path meets only the vertex u and $[p]_G = 1$ so that for q we may choose the empty word 1, which labels the empty path $u \rightarrow u$ and $[p]_G = [1]_G$.

If e is not a loop edge, then no power e^n or e^{-n} for $n \geq 2$ forms a path; therefore, if $|A| = 1$ the only possibilities for p are $f(f^{-1}f)^n$ and $(ff^{-1})^{n+1}$ for some $n \geq 0$ and $f \in \{e, e^{-1}\}$. In these cases the claim is obvious.

In the following, we use the notation of the series (5.1) and denote the Cayley graphs of H_k and G by \mathcal{H}_k and \mathcal{G} , respectively. So, let $|A| = k + 1$ for some $k \geq 1$, and let $\mathcal{A} = \langle A \rangle = \langle p \rangle$ be the subgraph of \mathcal{E} spanned by A , which, by definition, is the same as the subgraph of \mathcal{E} spanned by the path p (which therefore is connected). Abusing notation, we denote the labelled version of \mathcal{A} inside \mathcal{X}_1 also by \mathcal{A} and let $\varphi_u: \mathcal{H}_k \rightarrow \mathcal{X}_1$ be the canonical morphism mapping $1 \in \mathcal{H}_k$ to u ; let $\mathcal{A}_k \subseteq \mathcal{H}_k$ be the cover of \mathcal{A} in \mathcal{H}_k with $1 \in \mathcal{A}_k$ (that is, the connected component of $\varphi_u^{-1}(\mathcal{A})$ which contains the vertex 1). The path p in \mathcal{E} , or, more precisely, the path $\pi_u^{\mathcal{X}_1}(p)$ lifts to the path $\pi_1^{\mathcal{A}_k}(p)$. In particular, in \mathcal{A}_k there is a p -labelled path starting at 1. We consider the full $\leq A$ -coset extension $\text{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)$ and note that $\text{CE}(H_k, \mathcal{A}_k; B)$ is a subgraph of it. We also have the path $\pi_1^{\mathcal{G}}(p)$ in \mathcal{G} starting at 1 and being labelled p . The canonical morphism $\psi: \mathcal{G} \rightarrow \text{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)$ (mapping $1 \in \mathcal{G}$ to $1 \in \mathcal{A}_k$) maps $\pi_1^{\mathcal{G}}(p)$ to $\pi_1^{\text{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)}(p)$, but this path runs entirely in \mathcal{A}_k hence coincides with the path $\pi_1^{\mathcal{A}_k}(p)$ mentioned earlier.

By assumption, $[p]_G = [r]_G$ for some word $r \in \tilde{B}^*$. The paths $\pi_1^{\mathcal{G}}(p)$ and $\pi_1^{\mathcal{G}}(r)$ have the same terminal vertex, namely $[p]_G = [r]_G$. The path $\pi_1^{\mathcal{G}}(r)$ is mapped by ψ onto the path $\pi_1^{\text{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)}(r)$. But the B -component of 1 in $\text{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)$ is the full B -coset $1\mathcal{H}_k[B]$, which is contained in $\text{CE}(H_k, \mathcal{A}_k; B)$. So the latter graph contains a path labelled r starting at 1,

and that path $\pi_1^{\text{CE}(H_k, \mathcal{A}_k; B)}(r)$ actually runs inside $1\mathcal{H}_k[B]$. Since the paths $\pi_1^{\mathcal{G}}(r)$ and $\pi_1^{\mathcal{G}}(p)$ have the same terminal vertex, so have the paths

$$\pi_1^{1\mathcal{H}[B]}(r) = \pi_1^{\text{CE}(H_k, \mathcal{A}_k; B)}(r) \text{ and } \pi_1^{\mathcal{A}_k}(p).$$

It follows that the terminal vertex v' of $\pi_1^{\mathcal{A}_k}(p)$ is in $\mathcal{A}_k \cap 1\mathcal{H}_k[B]$. But $\mathcal{A}_k \cap 1\mathcal{H}_k[B]$ is just the B -component of 1 in \mathcal{A}_k which is a connected B -graph. Altogether, there exists a path $\pi: 1 \rightarrow v'$ running in $\mathcal{A}_k \cap 1\mathcal{H}_k[B]$; let $q \in \tilde{B}^*$ be the label of that path. By construction, $[q]_{H_k} = [r]_{H_k}$, hence $[q]_G = [r]_G$ since the expansion $H_k \leftarrow G$ is k -stable, and therefore also $[q]_G = [p]_G$. Finally, the canonical morphism $\varphi_u: \mathcal{H}_k \rightarrow \mathcal{X}_1$ (restricted to $1\mathcal{H}_k[B]$) maps $\pi = \pi_1^{\mathcal{A}_k \cap 1\mathcal{H}_k}(q)$ to the path $\pi_u^{\mathcal{X}_1}(q)$ with initial vertex $u = \varphi_u(1)$ and terminal vertex $v = \varphi_u(v')$ and label q . If we ignore the labelling then the latter path is the sequence q of edges in \mathcal{E} which forms a path $u \rightarrow v$. Altogether, q forms a path $u \rightarrow v$ in \mathcal{E} . \square

This proof sheds some light on the rôles that the components of \mathcal{Z}_k play in the transition $H_k \rightsquigarrow G_{k+1}$. If there is a word p with $\text{co}(p) = A$ and $|A| = k + 1$ such that p forms a path $u \rightarrow v$ in \mathcal{E} , and some letter $a \in A$ does not belong to the $\overline{\text{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)}$ (for $\mathcal{A} = \langle A \rangle$ and $B = A \setminus \{a\}$) guarantees that the next group G_{k+1} avoids every relation $p = r$ for any $r \in \tilde{B}^*$ (compare Remark 3.19) unless there exists a word $q \in \tilde{B}^*$ such that $[p]_{H_k} = [q]_{H_k}$ and q forms a path $u \rightarrow v$ in \mathcal{E} . From this point of view, namely to avoid all relations that would obstruct Lemma 5.6, it would be sufficient to let \mathcal{Z}_k be comprised of all graphs $\text{CE}(H_k, \mathcal{A}_k; B)$ of the mentioned kind (after making them weakly complete by extending edges to 2-cycles whenever needed). However, when attempting this approach, namely letting \mathcal{Z}_k be comprised of just all graphs of the mentioned form, the authors failed to prove k -stability of the expansion $H_k \leftarrow G_{k+1}$, and it is not clear whether or not k -stability can be achieved by this procedure. Hence, except for the graphs $\text{CE}(H_k, \mathcal{A}_k; B)$, which appear as subgraphs of the full coset extensions $\text{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)$, all the machinery used to set up the graph \mathcal{Z}_k — (augmented) clusters, (augmented) full coset extensions, all of Section 4 — serves to achieve k -stability of the transition $H_k \rightsquigarrow G_{k+1}$.

If, in Lemma 5.6, $[p]_G = 1$ then necessarily $u = v$ since in this case the path $\pi_1^{\mathcal{G}}(p)$ is closed and the canonical morphism $\varphi_u: \mathcal{G} \rightarrow \mathcal{X}_1$ maps this path onto the closed path $\pi_u^{\mathcal{X}_1}(p)$. The path p in \mathcal{E} obtained by ignoring the labelling then clearly is also closed. Iterated application of Lemma 5.6 leads to the following; for the definition of a content function C the reader recall Definition 3.4.

Corollary 5.7. *Let $p \in \tilde{E}^*$ be a word which forms a path $u \rightarrow v$ in \mathcal{E} ; then there exists a word $q \in \tilde{E}^*$ which uses only letters (i.e. edges) from the content $C([p]_G)$ (and/or their inverses) such that $[p]_G = [q]_G$ and q forms a*

path $u \rightarrow v$ in \mathcal{E} . If $C([p]_G) = \emptyset$, then $u = v$ and q is the empty word. If $C([p]_G) \neq \emptyset$, then the graph $\langle C([p]_G) \rangle = \langle \text{co}(q) \rangle$ is connected and contains the vertices u and v .

5.4. Final remark: pointlike conjecture for inverse monoids versus F -inverse cover problem. What can we say about the gap between these two problems? As already mentioned, the truth of the pointlike conjecture for inverse monoids follows from Ash’s result on inevitable labellings of graphs. The construction in Ash’s paper is quite involved and the groups constructed there are not very well traceable. However, in [4] it has been shown that the expansion Q^{Ab_p} of an A -generated group Q witnesses the pointlike sets of the inverse monoid $M(Q)$ and therefore is able to verify the pointlike conjecture for inverse monoids. For any prime p , the so-called *universal p -expansion* Q^{Ab_p} of Q is the largest A -generated expansion R of Q with kernel of $R \twoheadrightarrow Q$ an elementary Abelian p -group. This expansion can be obtained by the construction in (2.1), except that the E -group G used there is replaced with the free E -generated Abelian group of exponent p (which is the $|E|$ -fold direct product of cyclic groups of order p), in fact a very transparent group. Sufficient to verify the pointlike conjecture is an E -group which reflects the structure of the Cayley graph \mathcal{Q} of Q in a very weak sense: the graph spanned by the content of a word over \tilde{E} which forms a path $u \rightarrow v$ requires only a *connected component* containing u and v . From this point of view, it seems to be justified to say that the gap between the pointlike problem for inverse monoids and the F -inverse cover problem is huge, indeed.

As already mentioned, Henckell and Rhodes considered Problem 1.1 as a “stronger form” of the pointlike conjecture for inverse monoids. On the other hand, in the last sentence of their paper they wrote: “We do not necessarily believe [the F -inverse cover problem] has an affirmative answer.” So, in contrast to what is often reported, Henckell and Rhodes did not really conjecture that every finite inverse monoid does admit a finite F -inverse cover, but rather seem to have been undecided about this question. In fact, they seem to have had some feeling that the F -inverse cover problem might be hard.

REFERENCES

- [1] J. Almeida and M. Delgado, Sur certains systèmes d’équations avec contraintes dans un groupe libre, *Portugal. Math.* **56** (1999), no. 4, 409–417.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] K. Auinger and M. B. Szendrei, On F -inverse covers of inverse monoids, *J. Pure Appl. Algebra* **204** (2006), 493–506.

- [6] K. Auinger and A. Bors, Formations of finite groups with the M. Hall property, *J. Comb. Algebra* **3** (2019), no. 2, 113–152.
- [7] 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.
- [8] B. Billhardt, B. Singha, W. Sommanee, P. Thamkaew, J. Tiammee, A note on dual prehomomorphisms from a group into the Margolis–Meakin expansion of a group, *Semigroup Forum* **101** (2020), no. 3, 534–546.
- [9] J. Bitterlich, Finite F -inverse covers do exist, preprint (2018), arXiv:1808.05161.
- [10] J. Bitterlich, Investigations into the Universal Algebra of Hypergraph Coverings and Applications. Dissertation, Technische Universität Darmstadt, 2019.
- [11] K. Henckell and J. Rhodes, The theorem of Knast, the $PG = BG$ and type-II conjectures, Monoids and semigroups with applications (Berkeley, CA, 1989), 453–463, World Sci. Publ., River Edge, NJ, 1991.
- [12] 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.
- [13] I. Kapovich and A. Myasnikov, Stallings foldings and subgroups of free groups, *J. Algebra* **248** (2002), no. 2, 608–668.
- [14] M. V. Lawson, *Inverse Semigroups. The Theory of Partial Symmetries*, World Scientific, River Edge, NJ, 1998.
- [15] 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.
- [16] M. Otto, Highly acyclic groups, hypergraph covers, and the guarded fragment, *J. ACM* **59** (2012), no. 1, Art. 5, 40 pp.
- [17] M. Otto, Groupoids, hypergraphs, and symmetries in finite models. (English summary) 2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2013), 53–62, IEEE Computer Soc., Los Alamitos, CA, 2013.
- [18] M. Otto, Finite groupoids, finite coverings and symmetries in finite structures, preprint (2014) arXiv:1404.4599.
- [19] M. Otto, Acyclicity in finite groups and groupoids, preprint (2022), arXiv:1806.08664v5.
- [20] M. Petrich, *Inverse Semigroups*, Wiley, New York, 1984.
- [21] L. Ribes and P. Zalesskiĭ, On the profinite topology on a free group, *Bull. London Math. Soc.* **25** (1993), no. 1, 37–43.
- [22] J. R. Stallings, Topology of finite graphs, *Invent. Math.* **71** (1983), no. 3, 551–565.
- [23] N. Szakács, On the graph condition regarding the F -inverse cover problem, *Semigroup Forum* **92** (2016) no. 3, 551–558.
- [24] N. Szakács and M. B. Szendrei, On F -inverse covers of finite-above inverse monoids, *J. Algebra* **452** (2016), 42–65.
- [25] M. B. Szendrei, Some open problems in the structure theory of regular semigroups, *Semigroup Forum* **64** (2002), no. 2, 213–223.
- [26] M. B. Szendrei, Problems on finite regular semigroups, Proceedings of the International Symposium on Semigroups and Applications, 37–45, Dept. Math. Univ. Kerala, Thiruvananthapuram, 2007.

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

E-mail address: karl.auinger@univie.ac.at

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

E-mail address: otto@mathematik.tu-darmstadt.de