FINITE APPROXIMATION OF FREE GROUPS I: THE F-INVERSE COVER PROBLEM

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 $\mathrm{C}(g)$ of edges (the *content* of g) needed to represent g as a word over $\mathrm{C}(g) \cup \mathrm{C}(g)^{-1}$. 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 $u \longrightarrow v$ in $\mathcal{E}$ then there exists a G-equivalent word q which also forms a path $u \longrightarrow v$ 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. As the free group generated by E obviously has these properties, 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 [16], 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 any 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 $g \in G$.) The questions raised by Henckell and Rhodes in [16] concerned the algorithmic recognisability of certain subsets of M and relations on M for a given finite monoid M. These subsets and relations are defined by use of the collection of all subdirect products $T \subseteq M \times G$ of M with arbitrary finite groups G.

2020 *Mathematics Subject Classification.* 05C25, 05E18, 20M18, 20B99, 20F05, 20F65.

Key words and phrases. approximation of free groups in finite groups, permutation group, action graph, Cayley graph, finite inverse monoid, finite F-inverse cover.

Shortly after, all stated conjectures and one of the two problems (about *liftable tuples*) were verified respectively solved by Ash in his celebrated paper [5]. Roughly speaking, Ash proved that in the situations mentioned, and even beyond those, the collection of all subdirect products $T \subseteq M \times G$ of M with finite groups G has the same "computational power" as a particularly chosen "canonical" subdirect product $\tau \subseteq M \times F$ of M with some free group F. This is a strong form of approximation in finite groups of the free group F. The algorithmic recognisability of the aforementioned subsets and relations of M is an immediate consequence. The importance of Ash's paper went beyond its immediate task as in the following years interesting and deep connections with the profinite topology of the free group [27] and model theory [17] have been revealed and studied [2, 3].

Yet the second problem stated, which was called in [16] a "stronger form of the pointlike conjecture for inverse monoids", was not solved in Ash's paper and has since then attracted considerable attention [19, 20, 32, 33, 7, 31, 30, 11]. It asked:

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

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

When reading the paper [7] by Szendrei and the first author, the second author understood that a result by the third author [23, 24] about the existence of certain finite groupoids can be used to given an affirmative answer to Problem 1.1. He presented this solution in his dissertation [13] and his paper [12]. Later, some flaws were discovered in [23, 24] which, however, have been fixed in the meantime [25]. 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 [25] 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.

Since an infinite F-inverse cover can be constructed for every inverse monoid M by use of a free group F, the task for finite M is, to replace F by a suitable finite group H. The group H needs to have a sufficiently high combinatorial complexity in order to "simulate" the required behaviour of the free group F with respect to the monoid M. Hence the task is to *approximate* the free group F sufficiently well by a finite group H. What this exactly means in the present context is one of the essentials of the paper.

The paper is organised as follows: Section 2 collects prerequisites from inverse monoids, graphs and a proof that the existence of certain finite 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 obtain the required group in a construction which intends to "reflect the geometry" of a given finite graph $\mathcal{E}$ and thereby prove the main result of the paper (Lemma 2.5).

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 ${}^{-1}\colon M \to M$ 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 [20, Chapter 1, Theorem 1], inverse monoids may as well be characterised as monoids of partial bijections on a set, closed under composition of partial mappings and inversion. Therefore, while groups model symmetries of mathematical structures, inverse monoids (or semigroups) model partial symmetries, that is, symmetries between substructures of mathematical structures.

From basic facts of universal algebra it follows that every inverse monoid M admits a least congruence such that the corresponding quotient structure is a group. This congruence is usually denoted σ and it can be characterised as the least 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). In particular, $e \leq 1$ for every idempotent e of M. If an inverse monoid M is represented as a monoid of partial bijections, then the idempotents of M are exactly the restrictions of the identity 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 least congruence for which the natural order on the quotient is the identity relation, and, likewise as the least 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 [26] and Lawson [20].

An inverse monoid S is *F-inverse* if every σ-class of S 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 [9] and the literature cited there. An F-inverse monoid S is an *F-inverse cover* of the inverse monoid M if there exists a surjective idempotent separating homomorphism $S \to M$. As mentioned in the introduction, it has been an outstanding open problem whether every finite inverse monoid M admits a finite F-inverse cover.

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 $\widetilde{X} := X \cup X^{-1}$. The mapping $x \mapsto x^{-1}$ is extended to an involution of $\widetilde{X}$ by setting $(x^{-1})^{-1} = x$, for all $x \in X$. We let $\widetilde{X}^*$ be the free monoid over $\widetilde{X}$, which, subject to $(x_1 \cdots x_n)^{-1} = x_n^{-1} \cdots x_1^{-1}$ (where $x_i \in \widetilde{X}$), is the *free involutory monoid* over X. The elements of $\widetilde{X}^*$ are called *words over* $\widetilde{X}$, and we let 1 denote the *empty word*. A word $w \in \widetilde{X}^*$ is *reduced* if it does not contain any factor of the form xx^{-1} for $x \in \widetilde{X}$. Repeated deletion of such factors in a word w, until no one is present any more, leads to *the reduced form* $\mathrm{red}(w)$ of w.

We fix a non-empty set A (called alphabet in this context). An inverse monoid M together with a (not necessarily injective) mapping $i_M : A \to M$ (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 $M \to M/\rho$. A *morphism ψ from the A-generated inverse monoid M to the A-generated inverse monoid N* is a homomorphism $M \to N$ respecting generators from A, that is, satisfying $i_N = \psi \circ i_M$. If it exists, such a morphism is unique and surjective and is called *canonical morphism*, denoted $\psi : M \twoheadrightarrow N$. On a more formal level, an A-generated inverse monoid is an algebraic structure of the form $(M; \cdot, ^{-1}, 1, A)$ where every symbol $a \in A$ is interpreted in M as a constant (that is, as a nullary operation) via the assignment function i_M. Canonical morphisms of A-generated inverse monoids then are just homomorphisms of algebraic structures in the signature $\{\cdot, ^{-1}, 1\} \cup A$. If $M \twoheadrightarrow N$ then M is an *expansion of* N. In our usage, the term "expansion" just concerns the relationship between two individual A-generated inverse monoids M and N. This somehow deviates from the widespread use of that term standing for a functor on certain categories of monoids. The 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 least group congruence σ, is an A-generated group.

The assignment function i_M is usually not explicitly mentioned; it uniquely extends to a homomorphism $[\ \ ]_M\colon \widetilde{A}^* \to M$ (of involutory monoids). For every word $p \in \widetilde{A}^*$, $[p]_M$ is the *value of* p *in* M or simply the M*-value of* p. For two words $p, q \in \widetilde{A}^*$, 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 $[p]_M \neq [q]_M$.

Using the concept of "A-generatedness" we see that every inverse monoid admits an F-inverse cover. Indeed, let M be an inverse monoid; choose a set A with assignment function $i_M\colon A \to M$ so that M becomes A-generated and let F be the free A-generated group. Then the subdirect product

$$S := \{([w]_F, [w]_M)\colon w \in \widetilde{A}^*\} \subseteq F \times M \tag{2.1}$$

is an F-inverse cover of M. This is well known and it is easy to see. Indeed, the congruence σ on S can be described by

$$([u]_F, [u]_M)\ \sigma\ ([v]_F, [v]_M) \text{ if and only if } [u]_F = [v]_F.$$

Furthermore, for two words $u, v \in \widetilde{A}^*$ for which u is obtained from v by (successive) deletion of some factors of the form aa^{-1} ($a \in \widetilde{A}$) we have $[v]_M \leq [u]_M$ in any A-generated inverse monoid M. Consequently, for a given word $w \in \widetilde{A}^*$ the maximum element of the σ-class of $([w]_F, [w]_M)$ is the element $([\mathrm{red}(w)]_F, [\mathrm{red}(w)]_M) = ([w]_F, [\mathrm{red}(w)]_M)$. However, the inverse monoid S is infinite, no matter what M is. The Henckell–Rhodes problem then asks if in case of a finite inverse monoid M the infinite free group F in (2.1) may be replaced by some finite (A-generated) group H serving the same purpose. An affirmative answer to this question will be established in Theorem 2.4.

2.3. **Graphs.** In this paper, we consider the Serre definition [28] 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 *multidigraphs, directed multigraphs* or *quivers*. The following formalisation is convenient for our purposes. A *graph* $\mathcal{E}$ is a *two-sorted structure* $(V, K; \alpha, \omega, ^{-1})$ with V its set of *vertices*, K its set of *edges* (disjoint from V), with *incidence functions* $\alpha\colon K \to V$ and $\omega\colon K \to V$, selecting, for each edge e the *initial* vertex αe and the *terminal* vertex ωe, and *involution* $^{-1}\colon K \to K$ 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 $e\colon \underset{u}{\bullet}\longrightarrow\underset{v}{\bullet}$ and its inverse $e^{-1}\colon \underset{u}{\bullet}\longleftarrow\underset{v}{\bullet}$ as "the same edge but traversed in the opposite direction". A graph $(V, K; \alpha, \omega, ^{-1})$ is *oriented* if the edge set K is partitioned as $K = E \cup E^{-1} = \widetilde{E}$ 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 $\mathcal{E} = (V, \widetilde{E}; \alpha, \omega, {}^{-1})$.

A *subgraph* of the graph $\mathcal{E}$ is a substructure that is induced over a pair (V', K') of subsets $V' \subseteq V$ and $K' \subseteq K$ both of which are closed under the operations α and ${}^{-1}$ (and therefore also under ω). In particular, every pair (S, T) of subsets $S \subseteq V$ and $T \subseteq K$ generates a unique minimal subgraph $\langle (S, T) \rangle$ of $\mathcal{E}$ containing (S, T), which is the *subgraph of* $\mathcal{E}$ *spanned by* (S, T). An *automorphism* of a graph $\mathcal{E} = (V, K; \alpha, \omega, {}^{-1})$ is a pair of maps $\varphi = (\varphi_V, \varphi_K)$ with $\varphi_V \colon V \to V$, $\varphi_K \colon K \to K$ being bijections satisfying for all $e \in K$:

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

We note that the second equality is a consequence of the first and third. In the oriented case we require in addition that $\varphi_{\widetilde{E}}(E) = E$ and (therefore also) $\varphi_{\widetilde{E}}(E^{-1}) = E^{-1}$. 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 $\mathcal{E} = (V, K; \alpha, \omega, {}^{-1})$ *by the alphabet* A (an A-*labelling*, for short) is a mapping $\ell \colon K \to \widetilde{A}$ respecting the involution: $\ell(e^{-1}) = \ell(e)^{-1}$ for all $e \in K$. The labelling $\ell \colon K \to \widetilde{A}$ gives rise to an orientation of $\mathcal{E}$: setting $E := \{e \in K \colon \ell(e) \in A\}$ (positive edges) and $E^{-1} := \{e \in K \colon \ell(e) \in A^{-1}\}$ (negative edges), it follows that $E \cap E^{-1} = \varnothing$ and we get $K = \widetilde{E}$.

We consider A-labelled graphs as structures $(V, K; \alpha, \omega, {}^{-1}, \ell, A)$ 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 $\mathcal{K} = (V, K; \alpha, \omega, {}^{-1}, \ell, A)$ and $\mathcal{L} = (W, L; \alpha, \omega, {}^{-1}, \ell, A)$ be A-labelled graphs. *A morphism* $\varphi \colon \mathcal{K} \to \mathcal{L}$ *of A-labelled graphs* is a pair of mappings $\varphi = (\varphi_1, \varphi_2)$, $\varphi_1 \colon V \to W$, $\varphi_2 \colon K \to L$, both compatible with the operations α and ${}^{-1}$ (and therefore also ω) as well as with the labelling. Throughout the paper, in the situation of a morphism $\varphi = (\varphi_1, \varphi_2) \colon \mathcal{K} \to \mathcal{L}$ we shall write, for every vertex v [every edge e] of $\mathcal{K}$, $\varphi(v)$ instead of $\varphi_1(v)$ [$\varphi(e)$ instead of $\varphi_2(e)$].

A *congruence* Θ *on the A-labelled graph* $\mathcal{K} = (V, K; \alpha, \omega, {}^{-1}, \ell, A)$ is a pair (Θ_V, Θ_K) with Θ_V an equivalence relation on V, Θ_K an equivalence relation on K, compatible with the operations α and ${}^{-1}$ (therefore also ω) and respecting ℓ, that is:

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

and

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

The definition of the quotient graph $\mathcal{K}/\Theta$ for a congruence Θ is obvious, and we have the usual Homomorphism Theorem. As for images under morphisms, the congruence class $v\Theta_V$ of a vertex v [the congruence class $e\Theta_K$ of an edge e] will be denoted by $v\Theta$ [by $e\Theta$].

A *non-empty path* π in $\mathcal{E}$ is a sequence $\pi = e_1e_2\cdots e_n$ $(n \geq 1)$ 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 $\pi = e_1\cdots e_n$ is the graph spanned by the edges of π, that is, by the pair $(\varnothing, \{e_1, \ldots, e_n\})$; 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 $\widetilde{A}^*$.

2.4. **Cayley graphs of A-generated groups.** 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_Q:

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

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

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

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

for all $g, h \in Q$ and $(g,a) \in Q \times \widetilde{A}$, where h is an element of the acting group Q, g a vertex of $\mathcal{Q}$ and (g,a) an edge of $\mathcal{Q}$.

2.5. ***F*-inverse covers.** For a given finite A-generated inverse monoid M we intend to construct a finite A-generated group H for which the A-generated subdirect product

$$S := \{([w]_H, [w]_M) \colon w \in \widetilde{A}^*\} \subseteq H \times M \tag{2.2}$$

is an F-inverse cover of M. We start with a finite A-generated group Q such that, for all $w \in \widetilde{A}^*$, $[w]_Q = 1_Q$ implies that $[w]_M$ is an idempotent of M. Such a group Q can be found by representing M as an inverse monoid of partial bijections on a finite set X and extending the partial mappings $[a]_M$ ($a \in A$) to total permutations $\hat{a}$ on X or on some finite superset $Y \supseteq X$ and taking $Q := \langle \hat{a} \colon a \in A \rangle$, the group generated by the permutations $\hat{a}$ ($a \in A$). In semigroup theoretic terms this just means that the A-generated subdirect product

$$\{([w]_Q, [w]_M) \colon w \in \widetilde{A}^*\} \subseteq Q \times M$$

is an *E-unitary* cover of M.

The following lemma will be crucial. It is well known and readers familiar with the Margolis–Meakin-expansion $M(Q)$ of a group Q [21] will recognise that this lemma essentially proves the universal property of $M(Q)$. We present a proof in order to keep the paper more self-contained; it is a modified version of the proof of Lemma 4.6 in [9]. We fix some notation: for an A-generated group Q with Cayley graph $\mathcal{Q}$, $q \in Q$ and a word $w \in \widetilde{A}^*$ let $\pi_q^{\mathcal{Q}}(w)$ [resp. $\pi_1^{\mathcal{Q}}(w)$] denote the path in $\mathcal{Q}$ labelled w that starts at q [resp. 1_Q].

Lemma 2.1. *Let M be an A-generated inverse monoid and Q be an A-generated group such that, for all $w \in \widetilde{A}^*$, $[w]_Q = 1_Q$ implies that $[w]_M$ is an idempotent of M. Then, for any words $u, v \in \widetilde{A}^*$ for which $[u]_Q = [v]_Q$ and $\langle \pi_1^{\mathcal{Q}}(u) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(v) \rangle$ the inequality $[u]_M \geq [v]_M$ holds in M.*

Proof. The proof is by induction on the length $|u|$ of the word u. If $|u| = 0$, that is, if $u = 1$ is the empty word, then $[u]_Q = [v]_Q$ implies that $[v]_Q = [u]_Q = 1_Q$, whence $[v]_M$ is an idempotent so that $[v]_M \leq 1_M = [u]_M$. Let $|u| = 1$, that is, $u = a$ is a letter in $\widetilde{A}$. The assumptions $\langle \pi_1^{\mathcal{Q}}(a) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(v) \rangle$ and $[a]_Q = [v]_Q$ imply that either (i) $v = v_1 a v_2$ with $[v_1]_Q = [v_2]_Q = 1_Q$, or (ii) $v = v_1 a^{-1} v_2$ and $[v_1]_Q = [a]_Q = [v_2]_Q$. In case (i), $[v_1]_M, [v_2]_M \leq 1_M$, whence $[v]_M = [v_1 a v_2]_M = [v_1]_M [a]_M [v_2]_M \leq [a]_M$. In case (ii), $[v_1 a^{-1}]_Q = 1_Q = [a^{-1} v_2]_Q$, whence $[v_1 a^{-1}]_M, [a^{-1} v_2]_M \leq 1_M$ so that

$$[v]_M = [v_1 a^{-1} a a^{-1} v_2]_M = [v_1 a^{-1}]_M [a]_M [a^{-1} v_2]_M \leq [a]_M.$$

So let $|u| > 1$ and let $[v]_Q = [u]_Q$ and $\langle \pi_1^{\mathcal{Q}}(u) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(v) \rangle$, and assume that the statement of the lemma is true for all words u' with $|u'| < |u|$ and arbitrary v. Choose some factorisation $u = u_1 u_2$ with $|u_1|, |u_2| < |u|$. Let $q := [u_1]_Q$; the assumption $\langle \pi_1^{\mathcal{Q}}(u) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(v) \rangle$ implies that q is a vertex of $\langle \pi_1^{\mathcal{Q}}(v) \rangle$, i.e. the path $\pi_1^{\mathcal{Q}}(v)$ meets the vertex q. Let $v = v_1 v_2$ be a corresponding factorisation. That is, the terminal vertex of $\pi_1^{\mathcal{Q}}(v_1)$ is q. Then $[v v^{-1} v_1]_Q = [v_1]_Q = [u_1]_Q = q$ and clearly $\langle \pi_1^{\mathcal{Q}}(u_1) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(v v^{-1} v_1) \rangle$, whence $[v v^{-1} v_1]_M \leq [u_1]_M$ by the inductive hypothesis. Similarly, $[v_2 v^{-1} v]_Q = [v_2]_Q = [u_2]_Q$ and

$$\langle \pi_q^{\mathcal{Q}}(u_2) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(u) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(v) \rangle = \langle \pi_q^{\mathcal{Q}}(v_2 v^{-1} v) \rangle,$$

whence also $\langle \pi_1^{\mathcal{Q}}(u_2)\rangle \subseteq \langle \pi_1^{\mathcal{Q}}(v_2v^{-1}v)\rangle$ (here we apply the automorphism $x \mapsto {}^{q^{-1}}x$ of $\mathcal{Q}$). The inductive assumption implies $[v_2v^{-1}v]_M \leq [u_2]_M$. Altogether,

$$[v]_M = [vv^{-1}v_1]_M[v_2v^{-1}v]_M \leq [u_1]_M[u_2]_M = [u]_M.$$

□

For a given finite A-generated inverse monoid M and a finite A-generated group Q as above, we now seek to provide a finite expansion H of Q, for which the subdirect product (2.2) is a finite F-inverse cover of M. First we isolate an important property of groups generated by an alphabet.

Definition 2.2 (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 $\mathrm{C}(g)$ of X such that g is represented as a product of elements of $\mathrm{C}(g)$ and their inverses.

We need to define one further property, which will be crucial towards the construction of the desired group H.

Definition 2.3 (group reflecting the structure of a Cayley graph)**.** Let Q be an A-generated group with Cayley graph $\mathcal{Q}$, where $E := Q \times A$ is 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 word $p \in \widetilde{E}^*$ which forms a path $g \longrightarrow h$ in $\mathcal{Q}$, the following hold:
   (a) if $\mathrm{C}([p]_G) = \varnothing$, that is if $[p]_G = 1$, then $g = h$,
   (b) if $\mathrm{C}([p]_G) \neq \varnothing$, that is if $[p]_G \neq 1$, then there exists a word $q \in \widetilde{E}^*$ which also forms a path $g \longrightarrow h$ in $\mathcal{Q}$ and such that $[p]_G = [q]_G$ and q uses only edges of the content $\mathrm{C}([p]_G)$ of $[p]_G$ (and their inverses). In particular, the content $\mathrm{C}([p]_G)$ 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 a finite E-generated group reflecting the structure of the Cayley graph $\mathcal{Q}$ of Q. The existence of such a group G is guaranteed by Lemma 2.5, whose proof will be completed in Section 5. 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 $G \rtimes Q$:

$$H := \langle ([(1_Q, a)]_G, [a]_Q) : a \in A \rangle \subseteq G \rtimes Q. \tag{2.3}$$

Readers familiar with the Margolis–Meakin-expansion $M(Q)$ [21] will notice that the group H, in a sense, approximates $M(Q)$. The type of construction used for the group H occurs frequently in (semi)group theory, see e.g. Elston [15] or Almeida [1, Section 10]; that it can be useful for the construction of F-inverse covers is discussed in [7]. For a word $p \in \widetilde{A}^*$, the value of p in H is

$$[p]_H = ([\pi_1^{\mathcal{Q}}(p)]_G, [p]_Q) \tag{2.4}$$

where, again, $\pi_1^{\mathcal{Q}}(p)$ is the unique path in $\mathcal{Q}$ starting at 1_Q and being labelled p, interpreted as a word over $\widetilde{E}$. This is easily seen by induction on the length $|p|$ of p. In particular, H is an expansion of Q with canonical morphism $([\pi_1^{\mathcal{Q}}(p)]_G, [p]_Q) \mapsto [p]_Q$. We can now formulate the main theorem of this section, and derive it essentially based on Lemma 2.5.

Theorem 2.4. *Let M be a finite A-generated inverse monoid, Q a finite A-generated group such that, for all $w \in \widetilde{A}^*$, $[w]_Q = 1_Q$ implies that $[w]_M$ is an idempotent of M. For $E = Q \times A$, let G be a finite E-generated group (with content function* C*) which reflects the structure of the Cayley graph $\mathcal{Q}$ of Q (Definition 2.3), and let H be the group defined by* (2.3)*. Then the subdirect product*

$$S := \{([w]_H, [w]_M) : w \in \widetilde{A}^*\}$$

is a finite F-inverse cover of M.

Proof. It is clear that S is finite. The natural order $\leq$ on S is given by $(g, m) \leq (h, n)$ if and only if $g = h$ and $m \leq n$. The canonical morphism $S \twoheadrightarrow M$, $(g, m) \mapsto m$ is idempotent separating. We need to show that every σ-class $(g, m)\sigma$ has a greatest element. Note that $(g, m)\sigma = \{(h, n) \in S : h = g\}$. Let $w \in \widetilde{A}^*$ be a word such that $[w]_H = g$. Then $[w]_H = ([\pi_1^{\mathcal{Q}}(w)]_G, [w]_Q)$. Since G reflects the Cayley graph $\mathcal{Q}$ of Q there exists a word $\pi \in \widetilde{E}^*$ such that $[\pi]_G = [\pi_1^{\mathcal{Q}}(w)]_G$, π contains only edges from the content $\mathrm{C}([\pi_1^{\mathcal{Q}}(w)]_G)$ (and their inverses) and π forms a path $1_Q \longrightarrow [w]_Q$ in $\mathcal{Q}$. The path π is induced by some word $u \in \widetilde{A}^*$, that is, $\pi = \pi_1^{\mathcal{Q}}(u)$. As the terminal vertex of $\pi_1^{\mathcal{Q}}(u)$ is $[u]_Q$ we have $[w]_Q = [u]_Q$, together with $[\pi_1^{\mathcal{Q}}(w)]_G = [\pi_1^{\mathcal{Q}}(u)]_G$ therefore also $g = [w]_H = [u]_H$. While π and u are not necessarily uniquely determined by g, we note that $\langle \pi_1^{\mathcal{Q}}(u) \rangle$ is the graph spanned by $\mathrm{C}([\pi_1^{\mathcal{Q}}(w)]_G)$ and therefore $\langle \pi_1^{\mathcal{Q}}(u) \rangle$ is uniquely determined by $g = ([\pi_1^{\mathcal{Q}}(w)]_G, [w]_Q)$. In addition, for *every* $v \in \widetilde{A}^*$ for which $[v]_H = g$ we have

$$\langle \pi_1^{\mathcal{Q}}(u) \rangle = \langle \mathrm{C}([\pi_1^{\mathcal{Q}}(w)]_G) \rangle = \langle \mathrm{C}([\pi_1^{\mathcal{Q}}(v)]_G) \rangle,$$

which implies $\langle \pi_1^{\mathcal{Q}}(u) \rangle \subseteq \langle \pi_1^{\mathcal{Q}}(v) \rangle$ for every such v. It follows from Lemma 2.1 that $[v]_M \leq [u]_M$ and therefore also $(g, [v]_M) \leq (g, [u]_M)$ for every such v. This just says that $(g, [u]_M)$ is indeed the greatest element of $(g, m)\sigma$. $\square$

The F-inverse cover S of M in Theorem 2.4 is A-generated *as an inverse monoid*. In a sense, this is a strong form of F-inverse cover, as the original definition does not require the cover to be A-generated as an inverse monoid. For any finite A-generated inverse monoid M and any finite A-generated group K expanding an A-generated group H as in Theorem 2.4, there exists an F-inverse cover $U \subseteq K \times M$ of M with maximum group quotient K, albeit not necessarily one that is A-generated as an inverse monoid.

This can be seen as follows. Let M and H be as in Theorem 2.4 and let the A-generated group K be an expansion of H. Then the subdirect product

$$U := \{([u]_K, [v]_M) \in K \times M : u, v \in \widetilde{A}^*, [u]_H = [v]_H\} \tag{2.5}$$

is an F-inverse cover of M with maximum group quotient K. Indeed, let $([u]_K, [v]_M) \in U$, then $[u]_H = [v]_H$. By the proof of Theorem 2.4 there exists $w \in \widetilde{A}^*$ with $[w]_H = [u]_H$ and such that $[w]_M \geq [x]_M$ for *every* $x \in \widetilde{A}^*$ for which $[x]_H = [w]_H = [u]_H$. Since $[w]_H = [u]_H$ we have $([u]_K, [w]_M) \in U$; in addition, $([u]_K, [w]_M) \geq ([u]_K, [x]_M)$ for all $x \in \widetilde{A}^*$ for which $[x]_H = [u]_H$. It follows that $([u]_K, [w]_M)$ is the greatest element of the σ-class $([u]_K, [v]_M)\sigma$. It should be emphasised that the F-inverse cover U of M cannot be guaranteed to be A-generated as an inverse monoid, and the A-generated subdirect product of M and K contained in U cannot be guaranteed to be an F-inverse monoid.

2.6. **Excursion: symmetries.** An alternative perspective on A-generated inverse monoids, maybe broader in a model-theoretic sense, would view them as two-sorted structures in which the A-labelling of the generators is integrated in an explicit fashion. This leads to the two-sorted structure

$$\mathfrak{M} := (M, A; \iota, \cdot, ^{-1}, 1),$$

where the set A of labels forms a second sort along with first sort M, and $\iota \colon A \to M$ encodes the explicit A-labelling as a function, so that $\iota(A)$ becomes a subset of M. The A-generated inverse monoid as a structure $(M; \cdot, ^{-1}, 1, A)$ — with the algebraic signature $\{\cdot, ^{-1}, 1\}$ enriched by the set A of constant symbols interpreted as generators of the inverse monoid $(M; \cdot, ^{-1}, 1)$ — is clearly rigid since fixing the generators fixes everything. The two-sorted structure $\mathfrak{M}$, on the other hand, allows us to analyse internal symmetries induced by permutations of A: these precisely are the automorphisms of the two-sorted structure $\mathfrak{M}$, namely pairs of compatible permutations of the sets M and A that commute with $\iota \colon A \to M$ and the algebraic operations on M. As $\iota(A) \subseteq M$ is a generating set in $(M; \cdot, ^{-1}, 1)$, compatibility with the algebraic operations now implies that any such automorphism is uniquely determined by its action on A.

As our construction of the F-inverse cover S of M according to Theorem 2.4 deals with A-generated inverse monoids, the question arises to which extent this construction may also be symmetry-preserving – in the sense of commuting with permutations of the generator set A that induce automorphisms of $\mathfrak{M}$. And indeed, our construction can be made fully symmetry-preserving overall. Theorem 2.4 rests on the chain

$$M \rightsquigarrow Q \rightsquigarrow H \rightsquigarrow S := H \times_A M,$$

where $H \times_A M$ denotes the A-generated subdirect product of the A-generated group H and the A-generated inverse monoid M. The step from the A-generated group Q to the A-generated group H, which is the technically challenging construction behind Lemma 2.5, is really based on the Cayley graph $\mathcal{Q}$ of Q; from this Cayley graph $\mathcal{Q}$ we first obtain a $(Q \times A)$-generated group G, which reflects the structure of $\mathcal{Q}$ in the sense of Definition 2.3, and finally H as the group defined in (2.3). The Cayley graph $\mathcal{Q}$, as a graph that is edge-labelled by A, lifts every symmetry of the A-generated group Q in a canonical manner. The symmetry-preserving nature of the core construction in the passage from this Cayley graph $\mathcal{Q}$ to G is captured in Proposition 5.5 and to be further discussed in Section 5. That H as defined in (2.3) then carries all shared symmetries of G and $\mathcal{Q}$ is straightforward, and similarly for $S = H \times_A M$ in relation to H and M. In fact, all of these steps between Q and S, including the core passage from $\mathcal{Q}$ to G, can be seen to be symmetry-preserving (in terms of two-sorted presentations with explicit domains for the generator or label sets), simply because they can all be cast as explicit definitions (in more model-theoretic terms: as interpretations, albeit of a higher-order nature) of the target structures over the input structures, which cannot possibly violate isomorphism invariance.

Therefore the one step that curiously risks breaking symmetries, lies in the passage from M to Q. The straightforward recipe indicated above is to obtain a representation of M as an inverse monoid of partial bijections over some set X, which can then be extended to permutations of X to yield a permutation group Q that relates to M as required for Lemma 2.1. But this method crucially involves free choices of a representation over a suitable set X and of extensions from partial to global bijections over that X; and these choices can break symmetries.[1] This seeming obstacle can be overcome though, if we use for instance the canonical representation of the A-generated inverse monoid M by partial bijections over the set M itself, and then consider all possible extensions of the partial mappings $[a]_M \colon M \to M$ in parallel: for every possible extension of the family $([a]_M)_{a \in A}$ by total permutations $[\hat{a}]_M \supseteq [a]_M$ of M, we may use, for the set X, the disjoint union of copies of M, one for each choice of extensions, and the bijection induced by the instances of $[\hat{a}]_M$ in each copy of M as the generator set for Q in the symmetric group over X. This is but one of several variants for the step from M to Q that are explicitly definable over $\mathfrak{M}$ and therefore symmetry-preserving.

[1]E.g. two partial bijections that are identical over X may be extended to permutations of different orders.

2.7. **Excursion: pointlike conjecture for inverse monoids versus F-inverse cover problem.** What can we say about the gap between these two problems? Recall that for an A-generated inverse monoid M, an A-generated group H is a witness for the pointlike pairs of M if

$$\forall\, u_1, u_2 \in \widetilde{A}^* \;\exists\, v \in \widetilde{A}^* \colon [u_1]_H \neq [u_2]_H \;\text{ or }\; [u_1]_M, [u_2]_M \leq [v]_M. \tag{2.6}$$

On the other hand, the A-generated subdirect product

$$\{([w]_H, [w]_M) \colon w \in \widetilde{A}^*\} \subseteq H \times M$$

is an F-inverse cover of M provided that

$$\forall\, u_1, u_2 \in \widetilde{A}^* \;\exists\, v \in \widetilde{A}^* \colon [u_1]_H \neq [u_2]_H \text{ or } \begin{cases} [u_1]_M, [u_2]_M \leq [v]_M \\ \text{and} \\ [u_1]_H = [v]_H = [u_2]_H. \end{cases} \tag{2.7}$$

As shown in [6], the expansion $H = Q^{\mathbf{Ab}_p}$ of an A-generated group Q witnesses the pointlike sets of the inverse monoid M (Q in relation to M as in Lemma 2.1) and therefore the pointlike conjecture for inverse monoids is verified; in particular $H = Q^{\mathbf{Ab}_p}$ satisfies condition (2.6). For any prime p, the so-called *universal p-expansion* $Q^{\mathbf{Ab}_p}$ of Q is the largest A-generated expansion $R \twoheadrightarrow Q$ whose kernel is an elementary Abelian p-group. This expansion can be obtained by the construction in (2.3), 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 for the verification of 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 $\widetilde{E}$ which forms a path $u \longrightarrow v$ requires only a *connected component* containing u and v. The enormous effort we require in the remainder of the paper to construct an expansion H of Q that satisfies the seemingly innocent, additional condition $[u_1]_H = [v]_H = [u_2]_H$ in (2.7), indicates that the gap between the pointlike problem for inverse monoids and the F-inverse cover problem may indeed be huge.

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.

2.8. **The main result.** In order to prove Theorem 2.4 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 2.3. 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 2.5 (main lemma)**.** *For every finite connected oriented graph $\mathcal{E} = (V, \widetilde{E}; \alpha, \omega, ^{-1})$ 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 $p \in \widetilde{E}^*$) is closed under the deletion of generators and thus G has a content function* C *(Proposition 3.5).*
(3) *For any word $p \in \widetilde{E}^*$ which forms a path $u \longrightarrow v$ in $\mathcal{E}$ (with u and v not necessarily distinct vertices of $\mathcal{E}$) the following hold:*
   (a) *if $\mathrm{C}([p]_G) = \varnothing$ then $u = v$,*
   (b) *if $\mathrm{C}([p]_G) \neq \varnothing$ then there exists a word $q \in \widetilde{E}^*$ with $[p]_G = [q]_G$ such that q also forms a path $u \longrightarrow v$ in $\mathcal{E}$ and q only uses edges from the content $\mathrm{C}([p]_G)$ (and their inverses). In particular, $\mathrm{C}([p]_G)$ spans a connected subgraph of $\mathcal{E}$ containing u and v.*

Remark 2.6. The free group generated by E obviously enjoys properties (1)–(3) of Lemma 2.5. 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 [2, 3, 5, 17, 22], in contrast to [14] where such an approximation is not possible.

The remainder of the paper is devoted to proving Lemma 2.5. 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 those three sections are of a more general nature, may be of independent interest and will be of particular use in the follow-up paper [10].

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 be infinite in most cases, but must in some sense be regular (recognisable by a finite automaton). If necessary take a quotient $\mathcal{X}/{\equiv}$ of $\mathcal{X}$ which guarantees that the edge labels from A induce partial permutations on the vertex set. Finally form some *completion* $\overline{\mathcal{X}/{\equiv}}$ of $\mathcal{X}/{\equiv}$, through extending the partial permutations 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.

The simplest example of this procedure is the construction of a finite A-generated group which avoids a single relation $p = 1$ for a given reduced word $p \in \widetilde{A}^*$ — this provides a transparent and elegant proof that every free group is residually finite. A slightly more general application is the Biggs construction [4] providing a finite group that avoids all relations $p = 1$ for all reduced words p of length up to a given bound n — this has been used for the construction of finite regular graphs of large girth. A meanwhile classical and more advanced 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 [29]. Here a finite A-generated group is constructed that avoids all the (infinitely many) relations of the form $h = p$ where h runs through all elements of a finitely generated subgroup H of the free A-generated group F and p is a fixed element of $F \setminus H$. Many more examples can be found in [18, 8] and elsewhere. In his paper [5] Ash definitely developed some mastership of arguments of this kind. Independently, the third author has suggested a considerable refinement of this approach [22]. He proposed a construction which is inductive on the subsets of the generating set A in the sense that the kth 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 [22]. The constructions in this section and the results of the next section are of this flavour and are taken from the third author's [25].

3.1. **E-graphs and E-groups.** We slightly change perspective: since the edges of the graph $\mathcal{E}$ of Lemma 2.5 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 $a \in \widetilde{E}$, 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* [18] or *inverse automaton* [6, 8], to mention just two. In an E-graph $\mathcal{K}$, for every word $p \in \widetilde{E}^*$ 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 $p \in \widetilde{A}^*$ (for $A \subseteq E$) we write $u \xrightarrow{p} v$ and call π an *A-path* $u \longrightarrow v$; 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 $\widetilde{A}^*$. 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 $a \in \widetilde{E}$ exactly one edge f with initial vertex $\alpha f = u$ and label $\ell(f) = a$; in this case, for every word

$p \in \widetilde{E}^*$ 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 $p \in \widetilde{E}^*$, the mapping $[p] \colon V \to V$, $u \mapsto u \cdot p$ is a permutation of the vertex set V of $\mathcal{K}$. Thus the involutory monoid $\widetilde{E}^*$ acts on V by permutations on the right. The permutation group

$$\mathscr{T}(\mathcal{K}) := \{[p] \colon p \in \widetilde{E}^*\} \tag{3.1}$$

obtained this way, is called the *transition group* $\mathscr{T}(\mathcal{K})$ of the graph $\mathcal{K}$. This transition group $\mathscr{T}(\mathcal{K})$ is an E-generated group (E*-group* for short) in a natural way, the letter $e \in \widetilde{E}$ 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 $\mathscr{T}(\mathcal{K})$). Moreover, it may happen that distinct letters $e \neq f \in \widetilde{E}$ induce the same permutation.

A crucial fact concerning the transition group $G = \mathscr{T}(\mathcal{K})$ 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 $\varphi_u \colon \mathcal{G} \twoheadrightarrow \mathcal{C}$ from the Cayley graph $\mathcal{G}$ of G onto $\mathcal{C}$ for which $\varphi_u(1) = u$; we call φ_u the *canonical morphism* $\mathcal{G} \twoheadrightarrow \mathcal{C}$ *with respect to* u; occasionally we shall leave the vertex u undetermined and shall speak of *some* canonical morphism $\mathcal{G} \twoheadrightarrow \mathcal{C}$. The existence of these canonical morphisms will be frequently assumed without further mention. 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 $\varphi \colon \mathcal{G} \twoheadrightarrow \mathcal{C}$.

An E-graph $(V, K; \alpha, \omega, {}^{-1}, \ell, E)$ is *weakly complete* if, for every letter $a \in \widetilde{E}$, the partial permutation 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 $a \in \widetilde{E}$, 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 $p \in \widetilde{E}^*$ let $p_{a\to 1}$ 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: [2]

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

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

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

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

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 $\psi_A \colon G \twoheadrightarrow G[A]$ 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 $\mathcal{G} \twoheadrightarrow \overline{\mathcal{G}[A]}$, that is, $\mathcal{G}$ covers $\overline{\mathcal{G}[A]}$.

Suppose conversely that for every $A \subseteq E$ there is a canonical graph morphism $\mathcal{G} \twoheadrightarrow \overline{\mathcal{G}[A]}$. 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 $\mathcal{G} \twoheadrightarrow \overline{\mathcal{G}[B]}$ 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 $\pi_1^{\overline{\mathcal{G}[B]}}(p_{a\to 1})$ traverse the same edges except loop edges labelled $a^{\pm 1}$, and therefore visit the same vertices. So $\pi_1^{\overline{\mathcal{G}[B]}}(p_{a\to 1})$ is also closed, and as it runs entirely in $\mathcal{G}[B]$, it follows that $[p_{a\to 1}]_{G[B]} = 1$ and therefore $[p_{a\to 1}]_G = 1$. □

For a word $p \in \widetilde{E}^*$ the *content* $\mathrm{co}(p)$ 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 2.2). Indeed, assume that G is retractable. Then, for $p, q \in \widetilde{E}^*$ and $a \in E$ the equality $[p]_G = [q]_G$ implies $[p_{a\to 1}]_G = [q_{a\to 1}]_G$. Suppose now that $a \in \mathrm{co}(p)$ but $a \notin \mathrm{co}(q)$. Then the words q and $q_{a\to 1}$ are identical. Hence $[p]_G = [q]_G$ implies

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

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* $\mathrm{C}(g)$ of g is

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

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

The terminology is justified as $\mathrm{C} \colon g \mapsto \mathrm{C}(g)$ clearly is a content function in the sense of Definition 2.2. So we have shown the following.

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

$$G[A] \cap G[B] = G[A \cap B]. \tag{3.2}$$

Groups satisfying this condition for all $A, B \subseteq E$ have been called 2-*acyclic* by the third author in [22, 25]: 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 $u \longrightarrow v$; 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 [22, 25], 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:

$$\begin{aligned} & gG[A] = hG[A],\ hG[B] = kG[B] \text{ and } \ kG[C] = gG[C] \\ \Longrightarrow \quad & hG[A \cap B] \cap kG[B \cap C] \cap gG[C \cap A] \neq \varnothing, \end{aligned} \tag{3.3}$$

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

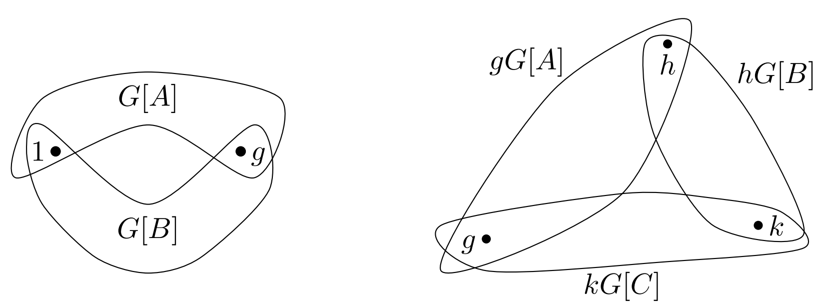

FIGURE 1

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).

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

$$H := \mathscr{T}\left(\mathcal{X} \sqcup \bigsqcup \{\overline{\mathcal{G}[C]} \colon C \subseteq E, |C| = k\}\right)$$

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 $p \in \widetilde{E}^*$ be a word with $|\mathrm{co}(p)| \leq k$ 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 $p \in \widetilde{A}^*$. So, let $B \subseteq E$ with $|B| = k$ and suppose that $p \notin \widetilde{B}^*$, 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 $\mathrm{co}(p) \setminus B$. 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 $\mathrm{co}(p) \setminus B$ — 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, which then also applies 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 $\mathcal{K}[A] \twoheadrightarrow \overline{\mathcal{K}[B]}^A$, where $\overline{\mathcal{K}[B]}^A$ denotes the trivial completion of $\mathcal{K}[B]$ as an A-graph, that is, loop edges labelled by letters form $A \setminus B$ (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 $\mathcal{K} \twoheadrightarrow \overline{\mathcal{G}[B]}$. Indeed, if $|B| = k$ then by definition of K and H, $\mathcal{K} \twoheadrightarrow \mathcal{H} \twoheadrightarrow \overline{\mathcal{G}[B]}$ since $\overline{\mathcal{G}[B]}$ is a component in the graph defining H as a transition group. If $|B| < k$ we may choose a set C with $B \subseteq C \subseteq A$ and $|C| = k$. Again there is a canonical morphism $\mathcal{K} \twoheadrightarrow \overline{\mathcal{G}[C]}$. Since $G[C]$ is retractable there is a canonical morphism $\mathcal{G}[C] \twoheadrightarrow \overline{\mathcal{G}[B]}^C$, where $\overline{\mathcal{G}[B]}^C$ is the trivial completion of $\mathcal{G}[B]$ as a C-graph. By adding loop edges for all labels not in C to all vertices, the morphism $\mathcal{G}[C] \twoheadrightarrow \overline{\mathcal{G}[B]}^C$ can be extended to a morphism $\overline{\mathcal{G}[C]} \twoheadrightarrow \overline{\mathcal{G}[B]}$. Composition with the morphism $\mathcal{K} \twoheadrightarrow \overline{\mathcal{G}[C]}$ then yields the desired 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 $\mathcal{K} \twoheadrightarrow \overline{\mathcal{G}[B]} \cong \overline{\mathcal{K}[B]}$ to $\mathcal{K}[A]$ provides the required morphism. □

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

$$G_1 \twoheadleftarrow G_2 \twoheadleftarrow \cdots \twoheadleftarrow G_{|E|} =: G \tag{3.4}$$

defined by an ascending sequence $\mathcal{X}_1 \subseteq \mathcal{X}_2 \subseteq \cdots \subseteq \mathcal{X}_{|E|}$ of complete E-graphs such that $G_k = \mathscr{T}(\mathcal{X}_k)$ is k-retractable and $G_{k+1} \twoheadrightarrow G_k$ 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 $p \in \widetilde{E}^*$ on $k+1$ letters which forms a path $u \xrightarrow{p} v$ in $\mathcal{E}$ and every letter $a \in A := \mathrm{co}(p)$ 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 $p = p_{a\to 1}$, 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 $G_{k+1} \twoheadrightarrow G_k$ is k-stable, the new components of $\mathcal{X}_{k+1}$ are constructed in a way so that their B-components for k-element subsets B of E have already occurred as subgraphs of $\mathcal{X}_k$. This turns out to be a challenging task. It crucially involves E-graphs whose A-components for $(k+1)$-element subsets A are designed so that their transition groups avoid certain new relations over A but preserve all relations over B for every $B \subsetneq A$. The latter is guaranteed, as already mentioned, by the fact that all B-components of the new components in $\mathcal{X}_{k+1}$ have been encountered already as subgraphs at earlier stages of the construction.

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.

In Sections 3.3.1 and 3.3.2 we discuss families of *clusters* and *coset extensions* whose A-components, as subgraphs of the E-graphs $\mathcal{X}_k$, provide the essential information for the setup of the expansions in the series (3.4); as discussed above, we need to account for their B-components for $B \subsetneq A \subseteq E$.

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

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

is an A-*cluster*. Note that $\mathsf{CL}(G[A], \mathbb{P})$ 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 $\mathsf{CL}(G[A], \mathbb{P})$. 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 $\mathsf{CL}(G[A], \mathbb{P})$ 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 $\mathcal{G}[B \cap C]$ (recall that retractability of $G[A]$ implies that $G[B \cap C] = G[B] \cap G[C]$). More precisely, let $\varphi \colon \bigsqcup_{B \in \mathbb{P}} \mathcal{G}[B] \to \mathcal{G}$ 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

$$\mathsf{CL}(G[A], \mathbb{P}) \cong \mathrm{im}(\varphi) \cong \bigsqcup_{B \in \mathbb{P}} \mathcal{G}[B] \,/\, \Theta. \tag{3.5}$$

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 $\mathsf{CL}(G[A], \mathbb{P})$ and $\mathsf{CL}(H[A], \mathbb{P})$ are isomorphic.*

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

We next analyse the structure of B-components of A-clusters for $B \subsetneq A$. Let $\mathbb{P} = \{A_1, \ldots, A_k\}$ 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

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

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) $v_i \in \mathcal{G}[A_i] \cap v\mathcal{G}[B]$. In order to describe the structure of the intersection $\mathsf{CL}(G[A], \mathbb{P}) \cap v\mathcal{G}[B]$ of a cluster $\mathsf{CL}(G[A], \mathbb{P})$ with a coset graph $v\mathcal{G}[B]$ we may ignore the constituent cosets $\mathcal{G}[A_i]$ of $\mathsf{CL}(G[A], \mathbb{P})$ having empty intersection with $v\mathcal{G}[B]$. Hence we may assume that $\mathcal{G}[A_i] \cap v\mathcal{G}[B] \neq \varnothing$ for every i.

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

$$\mathcal{G}[A_1] \cap \cdots \cap \mathcal{G}[A_k] \cap v\mathcal{G}[B] \neq \varnothing.$$

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

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

In order to simplify the notation we set $C := A_1 \cap \cdots \cap A_{t-1}$ and $D := A_t$; then $\bigcap_{i=1}^{t-1} \mathcal{G}[A_i] = \mathcal{G}[C]$. The situation is depicted in Figure 2 and the reader is invited to consult this illustration for the following argument. We need to exhibit an element in $\mathcal{G}[C \cap D] \cap v\mathcal{G}[B]$. So, let $u \in \big(\bigcap_{i=1}^{t-1} \mathcal{G}[A_i]\big) \cap v\mathcal{G}[B] = \mathcal{G}[C] \cap v\mathcal{G}[B]$ and $w \in \mathcal{G}[D] \cap v\mathcal{G}[B]$. Let $p \in \widetilde{C}^*$ be such that $[p]_G = u^{-1}$

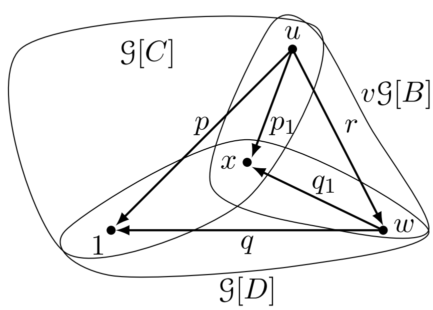

FIGURE 2

and $q \in \widetilde{D}^*$ be such that $[q]_G = w^{-1}$; moreover, let $r \in \widetilde{B}^*$ be a word which labels a path $u \longrightarrow w$ 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 $r \in \widetilde{B}^*$, we have $[p_1 q_1^{-1}]_G = [r]_G$, by retractability. Let $x := u \cdot p_1 = w \cdot q_1$. Then $p^{-1}p_1$ labels a path $1 \longrightarrow x$ and so does $q^{-1}q_1$. Since $p^{-1}p_1 \in \widetilde{C}^*$ and $q^{-1}q_1 \in \widetilde{D}^*$, it follows that $x \in \mathcal{G}[C] \cap \mathcal{G}[D] = \mathcal{G}[C \cap D]$. From $x = u \cdot p_1$ and $p_1 \in \widetilde{B}^*$ it follows that $x \in u\mathcal{G}[B] = v\mathcal{G}[B]$, altogether $x \in \mathcal{G}[C \cap D] \cap v\mathcal{G}[B]$. □

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, $1\mathcal{G}[C]$, $1\mathcal{G}[D]$ 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 $x \in G[A_1] \cap \cdots \cap G[A_k] \cap vG[B]$. Then $\mathcal{G}[A_i] = x^{-1}\mathcal{G}[A_i]$ for all i, and $x^{-1}v\mathcal{G}[B] = \mathcal{G}[B]$, so that

$$\begin{aligned} x^{-1}\big(\mathsf{CL}(G[A], \mathbb{P}) \cap v\mathcal{G}[B]\big) &= \textstyle\bigcup_{i=1}^k (\mathcal{G}[A_i] \cap \mathcal{G}[B]) \\ &= \textstyle\bigcup_{i=1}^k \mathcal{G}[A_i \cap B] \;=\; \mathsf{CL}(G[B], \mathbb{O}) \end{aligned} \tag{3.6}$$

where $\mathbb{O} = \{B \cap A_i \colon A_i \in \mathbb{P}\}$ (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 $\mathsf{CL}(G[A], \mathbb{P})$ is either a B-coset, that is, isomorphic with $\mathcal{G}[B]$, or isomorphic with the B-cluster $\mathsf{CL}(G[B], \mathbb{O})$ where $\mathbb{O} = \{C \cap B \colon C \in \mathbb{P}\}$ (some $C \cap B$ may be empty).*

Proof. The intersection $\mathsf{CL}(G[A], \mathbb{P}) \cap v\mathcal{G}[B]$ 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 $\mathsf{CL}(G[B], \mathbb{O})$, as indicated in (3.6). Now let v be a vertex of $\mathsf{CL}(G[A], \mathbb{P})$; then the B-component $\mathcal{B}$ of v in $\mathsf{CL}(G[A], \mathbb{P})$ is certainly contained in $\mathsf{CL}(G[A], \mathbb{P}) \cap v\mathcal{G}[B]$. Since the latter intersection is a B-cluster, it is connected and therefore $\mathcal{B}$ must coincide with this intersection. □

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. By Corollary 3.12, $\mathcal{B} = \mathsf{CL} \cap v\mathcal{G}[B]$ and $\mathcal{C} = \mathsf{CL} \cap w\mathcal{G}[C]$ 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. □

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

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

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* $\mathsf{CL}(G[A],\mathbb{P})$ *at* v. We have seen in Corollary 3.12 that the intersection $\mathsf{CL}(G[A],\mathbb{P}) \cap v\mathcal{G}[B]$ is a B-component of $\mathsf{CL}(G[A],\mathbb{P})$. It follows that the structure of the graph $\mathsf{CL}(G[A],\mathbb{P})\text{ⓥ}\mathcal{G}[B]$ only depends on the collection $\{G[C] \colon C \in \mathbb{P}\}$, the vertex v and $G[B]$ rather than on the entire group $G[A]$. Indeed, as can be seen in (3.5), the structure of $\mathsf{CL}(G[A],\mathbb{P})$ depends only on the graphs $\mathcal{G}[C]$ for $C \in \mathbb{P}$; furthermore, by Corollary 3.12, 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 $\iota \colon \mathcal{B} \to \mathcal{G}[B]$ determined by $v \mapsto 1$. (We may neglect the trivial case in Corollary 3.12, namely that $\mathcal{B} = \mathcal{G}[B]$: in that case, the augmented cluster would coincide with the original one.) The augmented cluster $\mathsf{CL}(G[A],\mathbb{P})\text{ⓥ}\mathcal{G}[B]$ can then be obtained as the disjoint union of $\mathsf{CL}(G[A],\mathbb{P})$ 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 $\mathsf{CL}(G[A],\mathbb{P})$ with $v := \varphi(u)$ there is an isomorphism of labelled graphs*

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

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. As noticed in Corollary 3.12, 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 $\mathsf{CL}(G[A],\mathbb{P})\text{ⓥ}\mathcal{G}[B]$ be a B-augmentation of the A-cluster $\mathsf{CL}(G[A],\mathbb{P})$ and let u be a vertex of this cluster. If the C-component $\mathcal{C}$ of u in $\mathsf{CL}(G[A],\mathbb{P})$ has empty intersection with the B-component $\mathcal{B}$ of v in $\mathsf{CL}(G[A],\mathbb{P})$ 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 $\mathcal{C} \cap \mathcal{B} \neq \varnothing$ 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 $\mathsf{CL}(G[A],\mathbb{P})\text{ⓥ}\mathcal{G}[B]$ 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, coset extensions, can be seen as a generalisation of clusters, but is more involved. It is a somewhat complex concept but it is perhaps *the* essential construction of the paper. Its definition will be developed over the next few pages and, in a sense, only culminates in the defining equation 3.10; but readers should bear in mind that its justification, including the non-trivial verification of its well-definedness, crucially relies on preparations expounded in the following pages.

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 $\widetilde{A}$ (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* $^{\subsetneq}$*A-coset extension* (with respect to G) if

$$\mathcal{B}_1 \cap \mathcal{B}_2 = \varnothing \text{ in } \mathcal{B} \Longrightarrow v_1\mathcal{G}[B_1] \cap v_2\mathcal{G}[B_2] = \varnothing \text{ in } v\mathcal{G}[B] \subseteq \mathcal{G}. \qquad (3.7)$$

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

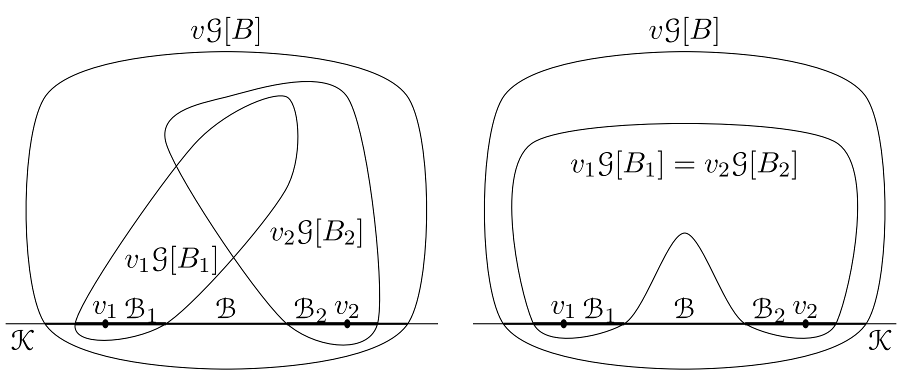

FIGURE 3

Definition 3.16 corresponds to the notion of *freeness* in [25], 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 ${}^{\subsetneq}A$-coset extension, then, for every $B \subseteq A$, every B-component $v\mathcal{K}[B]$ is admissible for ${}^{\subsetneq}B$-coset extension.

Now let $\mathcal{K}$ be a subgraph of $\mathcal{G}$ that is admissible for ${}^{\subsetneq}A$-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 $\mathsf{CE}(G, \mathcal{K}; B)$ be the graph obtained by extending each component $\mathcal{B}_i$ in $\mathcal{K}$ to the entire coset $v_i\mathcal{G}[B]$. So $\mathsf{CE}(G, \mathcal{K}; B)$ 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 $\mathsf{CE}(G, \mathcal{K}; B)$ 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 $\mathsf{CE}(G, \mathcal{K}; B)$. 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 $\iota_i \colon \mathcal{B}_i \to \mathcal{G}[B]$ be the unique graph monomorphism mapping v_i to 1. Then

$$\mathsf{CE}(G, \mathcal{K}; B) := \big(\mathcal{K} \cup \bigcup_{i=1}^{k} \mathcal{G}[B] \times \{i\}\big) \,/\, \Theta \tag{3.8}$$

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 (3.8) is a union of pairwise disjoint connected graphs and Θ is certainly a congruence relation. The resulting graph $\mathsf{CE}(G, \mathcal{K}; B)$ is the *B-coset extension of the A-graph* $\mathcal{K}$. The congruence Θ does not identify any two elements (edges or vertices) of $\mathcal{K}$ with each other, hence $\mathsf{CE}(G, \mathcal{K}; B)$ contains $\mathcal{K}$ as a subgraph in a canonical way which, in this context, is called the *skeleton* of $\mathsf{CE}(G, \mathcal{K}; B)$. For $v_i \in \mathcal{B}_i \subseteq \mathcal{K} \subseteq \mathsf{CE}(G, \mathcal{K}; B)$ the B-component of v_i in $\mathsf{CE}(G, \mathcal{K}; B)$ is isomorphic with the coset graph $\mathcal{G}[B]$. Hence these B-components of $\mathsf{CE}(G, \mathcal{K}; B)$ will also be denoted by $v_i\mathcal{G}[B]$ and addressed as *constituent cosets* of $\mathsf{CE}(G, \mathcal{K}; B)$ in this rôle.

For $C \subsetneq B \subsetneq A$, condition (3.7) of Definition 3.16 (by taking $B_1 = C = B_2$) implies that $\mathsf{CE}(G, \mathcal{K}; C)$ is realised as a subgraph of $\mathsf{CE}(G, \mathcal{K}; B)$. Moreover, for $C_1, C_2 \subsetneq B$, $C_1 \neq C_2$, once more condition (3.7) (this time taking $C_1 = B_1 \neq B_2 = C_2$) implies that

$$\mathsf{CE}(G, \mathcal{K}; C_1) \cap \mathsf{CE}(G, \mathcal{K}; C_2) = \mathsf{CE}(G, \mathcal{K}; C_1 \cap C_2) \tag{3.9}$$

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

$$\mathsf{CE}(G, \mathcal{K}; \mathbb{P}) := \Big(\bigcup \big\{\mathsf{CE}(G, \mathcal{K}; B) \times \{B\} \colon B \in \mathbb{P}\big\}\Big) \,/\, \Psi \tag{3.10}$$

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

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

In other words, an edge or a vertex of $\mathsf{CE}(G,\mathcal{K};B_1)$ is identified with one in $\mathsf{CE}(G,\mathcal{K};B_2)$ if they represent the same element in $\mathsf{CE}(G,\mathcal{K};B_1\cap B_2)$. Transitivity of Ψ follows from (3.9): indeed, for $i=1,2,3$, let $B_i\in\mathbb{P}$ and $x_i\in\mathsf{CE}(G,\mathcal{K};B_i)$ be such that $(x_1,B_1)\ \Psi\ (x_2,B_2)$ and $(x_2,B_2)\ \Psi\ (x_3,B_3)$. Then

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

so that

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

by application of (3.9) 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 $\mathsf{CE}(G,\mathcal{K};B_2)$. Provided that $B\in\mathbb{P}$, the coset extension $\mathsf{CE}(G,\mathcal{K};B)$ is embedded in $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$ 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 $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$ are the *constituent cosets* of $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$ and the subgraph $\mathcal{K}$ is the *skeleton* of $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$.

Geometrically, the coset extension $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$ can be viewed as follows. For every $B\in\mathbb{P}$ consider $\mathsf{CE}(G,\mathcal{K};B)$ 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 $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$ then is the union

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

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

$$\mathsf{CE}(G,\mathcal{K};B_1)\cap\mathsf{CE}(G,\mathcal{K};B_2) = \mathsf{CE}(G,\mathcal{K};B_1\cap B_2). \tag{3.11}$$

This is reminiscent of (3.9) 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 $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$ (rather than in $\mathsf{CE}(G,\mathcal{K};B)$). Moreover, condition (3.11) can be reformulated as a condition analogous to (3.7): for any $B_1,B_2\in\mathbb{P}$ and vertices $v_1,v_2\in\mathcal{K}$:

$$v_1\mathcal{K}[B_1]\cap v_2\mathcal{K}[B_2] = \varnothing \Longrightarrow v_1\mathcal{G}[B_1]\cap v_2\mathcal{G}[B_2] = \varnothing \tag{3.12}$$

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

If every label of $\mathcal{K}$ appears in some member B of $\mathbb{P}$, then $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$ is weakly complete since every edge of $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$ 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 $\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$ the *full* $\subsetneq A$-coset extension of $\mathcal{K}$. In case $\mathcal{K}=\{v\}$ (one vertex, no edge) the $\mathbb{P}$-coset extension $\mathsf{CE}(G,\mathcal{K};\mathbb{P})$ reduces to the cluster $\mathsf{CL}(G[A],\mathbb{P})$.

Remark 3.17. An A-graph $\mathcal{K}$ which is admissible for ${}^{\subsetneq}\!A$-coset extension may actually only contain edges labelled by letters (and their inverses) from some set $B \subsetneq A$. In this case $\mathsf{CE}(G, \mathcal{K}; B) \cong \mathcal{G}[B]$; however, this is not in conflict with the definition of the full ${}^{\subsetneq}\!A$-coset extension. For sets $C \subsetneq A$ with $C \nsubseteq B$, 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 $\mathsf{CE}(G, \mathcal{K}; C)$.

We continue with further investigations of ${}^{\subsetneq}\!A$-coset extensions.

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

Proof. We first establish a unique extension $\iota_B\colon \mathsf{CE}(G, \mathcal{K}; B) \to \mathcal{G}[A]$ for each $B \in \mathbb{P}$. Let $\mathcal{B}_1, \ldots, \mathcal{B}_k$ 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 $\kappa_i\colon \mathcal{G}[B] \times \{i\} \to \mathcal{G}[A]$ 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

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

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 $\iota_i\colon \mathcal{B}_i \to \mathcal{G}[B]$ is the unique graph monomorphism mapping v_i to 1 that occurs in the definition of $\mathsf{CE}(G, \mathcal{K}; B)$. It follows that the congruence Θ in (3.8) is contained in the kernel of κ and hence κ factors through $\mathsf{CE}(G, \mathcal{K}; B)$ as $\kappa = \iota_B \circ \pi_\Theta$ (where π_Θ is the canonical projection $\pi_\Theta(x) = x\Theta$).

Next consider the disjoint union

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

and let

$$\kappa_{\mathbb{P}} := \bigcup_{B \in \mathbb{P}} \boldsymbol{\iota}_B\colon \bigcup_{B \in \mathbb{P}} \mathsf{CE}(G, \mathcal{K}; B) \times \{B\} \to \mathcal{G}[A]$$

where $\boldsymbol{\iota}_B\colon \mathsf{CE}(G, \mathcal{K}; B) \times \{B\} \to \mathcal{G}[A]$ is defined by $\boldsymbol{\iota}_B(x, B) = \iota_B(x)$. Similar to Θ and κ, the congruence Ψ that occurs in (3.10) is contained in the kernel of $\kappa_{\mathbb{P}}$, whence $\kappa_{\mathbb{P}}$ factors through $\mathsf{CE}(G, \mathcal{K}; \mathbb{P})$ as $\kappa_{\mathbb{P}} = \iota_{\mathbb{P}} \circ \pi_\Psi$ for some unique morphism $\iota_{\mathbb{P}}\colon \mathsf{CE}(G, \mathcal{K}; \mathbb{P}) \to \mathcal{G}[A]$ (with π_Ψ being again the projection $x \mapsto x\Psi$). □

The morphism $\iota_B\colon \mathsf{CE}(G, \mathcal{K}; B) \to \mathcal{G}[A]$ 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 $\mathsf{CE}(G, \mathcal{K}; B)$. 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: $v_i\mathcal{G}[B] = v_j\mathcal{G}[B] =: v\mathcal{G}[B]$. In this case, ι_B maps $v_i\mathcal{G}[B]$ as well as $v_j\mathcal{G}[B]$ onto $v\mathcal{G}[B] \subseteq \mathcal{G}[A]$, although $\iota_B(v_i\mathcal{K}[B])$ and $\iota_B(v_j\mathcal{K}[B])$ are distinct (and hence disjoint) B-components of $\mathcal{K}$ within $v\mathcal{G}[B] \subseteq \mathcal{G}[A]$ (see Figure 4). 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 (some supergraph of) $\mathsf{CE}(G, \mathcal{K}; B)$. 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 $\widetilde{B}$-word, essentially because the graph $\mathsf{CE}(G, \mathcal{K}; B)$ unfolds the subgraph $\mathcal{K} \cup \bigcup v_i\mathcal{G}[B]$ of $\mathcal{G}[A]$ that arises as the image of $\mathsf{CE}(G, \mathcal{K}; B)$ under ι_B (see Figure 4).

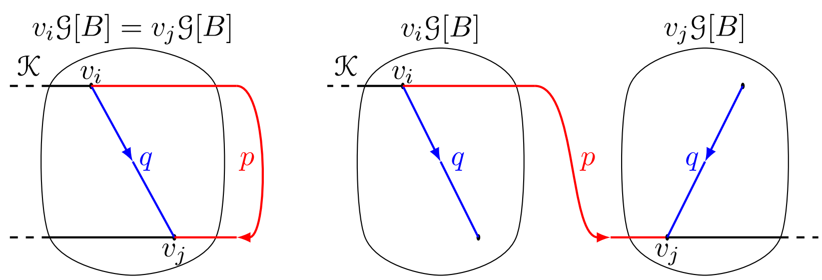

FIGURE 4. Part of $\mathcal{K} \cup \bigcup_{t=1}^{k} v_t\mathcal{G}[B] \subseteq \mathcal{G}[A]$ and of $\mathsf{CE}(G, \mathcal{K}; B)$

Let $\mathcal{K}$ be a connected A-graph admissible for ${}^{\subsetneq}\!A$-coset extension, let $B \subsetneq A$ and let $\mathcal{B} = v\mathcal{K}[B] \subseteq \mathcal{K}$ be the B-component of some vertex v in $\mathcal{K}$. By construction of $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$,

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

We are able to refine this chain as follows: $\mathcal{B}$ is itself admissible for ${}^{\subsetneq}\!B$-coset extension and hence $\mathsf{CE}(G, \mathcal{B}; \mathbb{P}_B)$ is well defined. Admissibility of $\mathcal{K}$ (Definition 3.16) implies that in this case the morphism $\iota_{\mathbb{P}_B} \colon \mathsf{CE}(G, \mathcal{B}; \mathbb{P}_B) \to \mathcal{G}[B]$ of Proposition 3.18 is injective. Indeed, $\iota_{\mathbb{P}_B}$ is injective on the skeleton $\mathcal{B}$, and on every constituent coset $v\mathcal{G}[C]$ for any $C \subsetneq B$ and any vertex v. If there were vertices $x \neq y$ such that $\iota_{\mathbb{P}_B}(x) = \iota_{\mathbb{P}_B}(y)$, then x and y would belong to two distinct constituent cosets $x \in v_1\mathcal{G}[B_1]$ and $y \in v_2\mathcal{G}[B_2]$ ($B_1, B_2 \subsetneq B$, possibly $B_1 = B_2$) so that x and y would coincide as elements of $v_1\mathcal{G}[B] = v_2\mathcal{G}[B]$. But this is excluded by Definition 3.16. Hence we get the following.

Lemma 3.20. *Let $\mathcal{K}$ be a subgraph of $\mathcal{G}[A]$ which is admissible for ${}^{\subsetneq}\!A$-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 ${}^{\subsetneq}B$-coset extension and the morphism $\iota_{\mathbb{P}_B} \colon \mathsf{CE}(G, \mathcal{B}; \mathbb{P}_B) \to \mathcal{G}[B]$ is injective. In particular, for any vertex $v \in \mathcal{B}$,*

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

Another consequence concerns connectivity in the graph $\mathcal{K}$; it will be of significant use later. In terms of [25] this means that a graph $\mathcal{K}$ which is admissible for ${}^{\subsetneq}A$-coset extension is 2-acyclic.

Lemma 3.21. *Suppose that the graph $\mathcal{K} \subseteq \mathcal{G}$ is admissible for ${}^{\subsetneq}A$-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 $\iota_{\mathbb{P}_A} \colon \mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A) \to \mathcal{G}[A]$. 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 ${}^{\subsetneq}A$-coset extension. Recall that the full ${}^{\subsetneq}A$-coset extension $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ can be seen as the union $\bigcup_{B \subsetneq A} \mathsf{CE}(G, \mathcal{K}, B)$ where for $B, C \subsetneq A$,

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

Every vertex x of $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ is sitting in some $\mathsf{CE}(G, \mathcal{K}; B)$, and, inside $\mathsf{CE}(G, \mathcal{K}; B)$ 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 $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$; 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 $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$; 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. It follows from (3.11) that every one-vertex subgraph of $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ has unique minimal support.

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

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

(1) every B-component $\mathcal{B}$ of $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ 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 [25].

Proposition 3.23. *Suppose that $\mathcal{K} \subseteq \mathcal{G}[A]$ is admissible for $^{\subsetneq}A$-coset extension and that the full $^{\subsetneq}A$-coset extension $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ 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 $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$.*

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, $\mathcal{B} \cong \mathsf{CL}(G[B], \{B_1, \ldots, B_k\})$ and $\mathcal{C} \cong \mathsf{CL}(G[C], \{C_1, \ldots, C_l\})$ 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

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

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 $v_1\mathcal{G}[F_1]\cap v_2\mathcal{G}[F_2] \neq \varnothing$, which by (3.12) implies $v_1\mathcal{K}[F_1]\cap v_2\mathcal{K}[F_2] \neq \varnothing$. 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 $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ and (3.11) 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 $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ formed by the union $v_1\mathcal{G}[F_1]\cup v_2\mathcal{G}[F_2]$ is isomorphic with the cluster $\mathsf{CL}(G[A], \{F_1, F_2\})$, 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

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

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 $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$; 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 $\mathcal{C} \cong \mathsf{CL}(G[C], \{C_1, \dots, C_l\})$ 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 $u_1 \longrightarrow y$ and also a $(B \cap C)$-path $y \longrightarrow u_2$, and altogether there is a $(B \cap C)$-path $u_1 \longrightarrow u_2$.

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 ${}^{\subsetneq}A$-coset extension and suppose that the full ${}^{\subsetneq}A$-coset extension $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ has the cluster property. For a vertex $v \in \mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ 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 $\iota\colon \mathcal{B} \hookrightarrow \mathcal{G}[B]$ (which is unique if one additionally

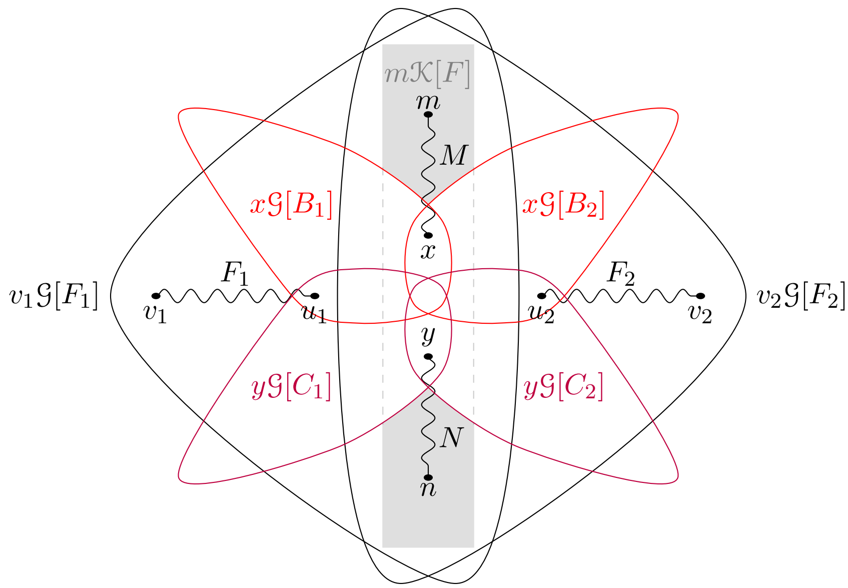

FIGURE 5

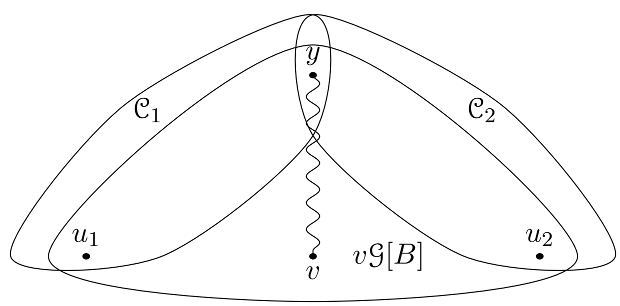

FIGURE 6

assumes that $\iota(v) = 1$). We define the *B-augmentation at v of* $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ by

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

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 $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A) \textcircled{v} \mathcal{G}[B]$ can be written as the union

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

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

Proposition 3.24. *Let $B, C \subsetneq A$ and $\mathcal{K}$ be admissible for $^{\subsetneq}A$-coset extension and such that the full $^{\subsetneq}A$-coset extension $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ enjoys the cluster property. Then every C-component of any B-augmented full coset* *extension* $\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)ⓥ\mathcal{G}[B]$ *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 $\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)ⓥ\mathcal{G}[B]$. If $\mathcal{C}\subseteq\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$ 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 $\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$ nor in $v\mathcal{G}[B]$. We have

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

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 $\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$). Let $\mathcal{B}_v$ be the B-component of v in $\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$. Our assumption implies that $\mathcal{C}\cap\mathcal{B}_v\neq\varnothing$. 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 $\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$, 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 $\mathcal{C}\subseteq\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$. 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

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

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

4. Two crucial inductive procedures

In this section we formulate and prove two important 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 (Definition 4.2 below). Assume, as above, that $|A|\geq 2$, that $G[A]$ is retractable and that $\mathcal{K}\subseteq\mathcal{G}[A]$ is admissible for ${}^{\subsetneq}A$-coset extension.

Definition 4.1 (embedded coset extension)**.** The full coset extension

$$\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$$

is *embedded* if the morphism $\iota_{\mathbb{P}_A}\colon\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)\to\mathcal{G}[A]$ (of Proposition 3.18) is an embedding.

Definition 4.2 (bridge freeness)**.** The embedded full coset extension

$$\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$$

is *bridge free* in $\mathcal{G}[A]$ if for every $B\subsetneq A$, if two vertices $u,v\in\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)\subseteq\mathcal{G}[A]$ (as per Definition 4.1) are B-connected in $\mathcal{G}[A]$, then they are B-connected even in $\mathsf{CE}(G,\mathcal{K};\mathbb{P}_A)$.

The two above-mentioned technical results will, in fact, be two inductive procedures — *forward induction* (Theorem 4.5) and *upward induction* (Theorem 4.7). Roughly speaking, forward induction guarantees that bridge freeness implies the cluster property — in the same group but with the number of letters being increased by one; upward induction, on the other hand, guarantees that the cluster property implies bridge freeness — with respect to the same set of letters but for the next group. For the construction of the series (3.4), these two procedures are applied alternatingly; the essence of the whole procedure is as follows (details will be worked out in Section 5.2). Suppose we have already defined the k-retractable group G_k. We apply Theorem 3.8 and produce a $(k+1)$-retractable and k-stable expansion H_k of G_k. Then take any connected A-subgraph $\mathcal{L}$ of the Cayley graph $\mathcal{H}_k$ of H_k for a subset $A \subseteq E$ of size $k+1$ and assume that $\mathcal{L}$ is admissible for ${}^{\subsetneq}\!A$-coset extension (with respect to H_k). For $B \subsetneq A$, all B-components $v\mathcal{L}[B]$ of $\mathcal{L}$ are subgraphs of $\mathcal{H}_k[B]$ and hence of $\mathcal{G}_k[B]$, by k-stability. Assuming inductively that all corresponding coset extensions $\mathsf{CE}(G_k, v\mathcal{L}[B]; \mathbb{P}_B)$ are bridge-free, the same is true for the corresponding coset extensions $\mathsf{CE}(H_k, v\mathcal{L}[B]; \mathbb{P}_B)$ with respect to H_k. *Forward induction* (Theorem 4.5) now implies that the coset extension $\mathsf{CE}(H_k, \mathcal{L}; \mathbb{P}_A)$ of the A-graph $\mathcal{L}$ has the cluster property. Finally, *upward induction* (Theorem 4.7) implies that for a suitable k-stable expansion G_{k+1} of H_k, any $\mathcal{G}_{k+1}$-cover $\hat{\mathcal{L}}$ of $\mathcal{L}$ is admissible for ${}^{\subsetneq}\!A$-coset extension (with respect to G_{k+1}) and that the coset extension $\mathsf{CE}(G_{k+1}, \hat{\mathcal{L}}; \mathbb{P}_A)$ is bridge-free (for a precise definition of *cover* see Definition 4.6 below).

The following lemma is the essential technical step to obtain the inductive procedure *forward induction* (Theorem 4.5). For this lemma 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 ${}^{\subsetneq}\!A$-coset extension, then all its B-components $v\mathcal{L}[B]$, for $B \subsetneq A$, are admissible for ${}^{\subsetneq}\!B$-coset extension and the morphisms of Proposition 3.18 are embeddings $\mathsf{CE}(H, v\mathcal{L}[B]; \mathbb{P}_B) \hookrightarrow v\mathcal{H}[B]$.

Lemma 4.3. *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 ${}^{\subsetneq}\!A$-coset extension. Assume that for all $B \subsetneq A$ and every vertex $v \in \mathcal{L}$, the full ${}^{\subsetneq}\!B$-coset extension $\mathsf{CE}(H, v\mathcal{L}[B]; \mathbb{P}_B)$*

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

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

Proof. Let $B \subsetneq A$, let $\mathcal{B}$ be a B-component of $\mathsf{CE}(H, \mathcal{L}; \mathbb{P}_A)$ 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 $\mathsf{CE}(H, \mathcal{L}; \mathbb{P}_A)$, 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 $\mathcal{B} \cap v_1\mathcal{H}[A_1] \neq \varnothing$, 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 $\mathsf{CE}(H, \mathcal{L}; \mathbb{P}_A)$ 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 ${}^{\subsetneq}A_i$-coset extension and the full ${}^{\subsetneq}A_i$-coset extension $\mathsf{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i})$ 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 $\mathcal{B}_i \cap \mathsf{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i}) \neq \varnothing$ — once more we take into account that

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

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

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

is connected. By assumption (2) therefore, $\mathcal{B}_i \cap \mathsf{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i})$ has unique minimal support in $\mathsf{CE}(H, \mathcal{A}_i; \mathbb{P}_{A_i})$, say (D_i, u_i). But then the pair (D_i, u_i) also provides unique minimal support of $\mathcal{B}_i$ in $\mathsf{CE}(H, \mathcal{L}; \mathbb{P}_A)$. 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 $\mathsf{CE}(H, \mathcal{L}; \mathbb{P}_A)$ 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 $s_i \longrightarrow u_i$ 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 $\mathcal{B}_i \cap \mathcal{B}_j \neq \varnothing$; after some renumbering we may assume that $\mathcal{B}_1 \cap \mathcal{B}_2 \neq \varnothing$. Then also $v_1\mathcal{H}[A_1] \cap v_2\mathcal{H}[A_2] \neq \varnothing$; from (3.12) we get $\mathcal{A}_1 \cap \mathcal{A}_2 \neq \varnothing$ 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 (3.11) 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 $\mathsf{CE}(H, \mathcal{A}_1 \cap \mathcal{A}_2; \mathbb{P}_{A_1 \cap A_2})$, (D, u) say, which (as for $\mathcal{B}_1$) provides unique minimal support of $\mathcal{B}_1 \cap \mathcal{B}_2$ in $\mathsf{CE}(H, \mathcal{L}; \mathbb{P}_A)$. 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

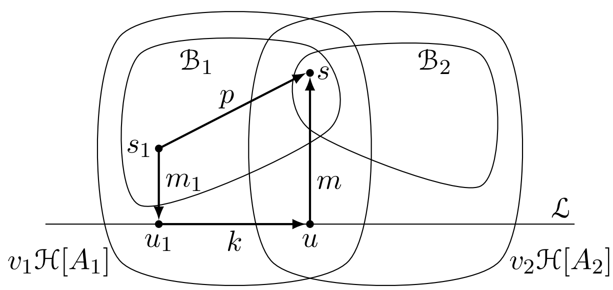

Figure 7

that $D \subseteq A_1 \cap A_2$ 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 $u_1 \longrightarrow u$ labelled k, say, which runs inside $\mathcal{A}_1$, and a D-path $u \longrightarrow s$ labelled m. Altogether, there is a D-path $s_1 \longrightarrow s$ 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 $s_1 \longrightarrow s$ 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 $s_1 \longrightarrow s$ 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 that $D \subseteq D_1$ 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 $D_2 = D$. 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 $s_1, \ldots, s_k \in \mathcal{B}_1 \cap \cdots \cap \mathcal{B}_k$ and all these $\mathcal{B}_i$ have

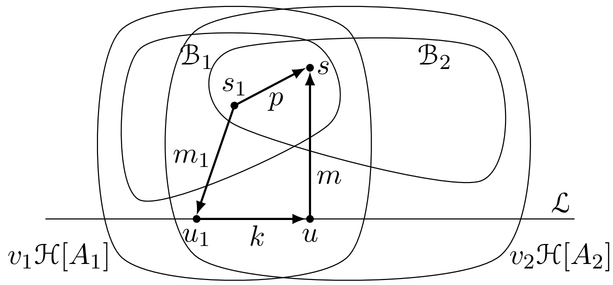

FIGURE 8

unique minimal support (D, u) attained at all these s_i. Again there are $j \in \{1, \ldots, k\}$ and $i \in \{k+1, \ldots, n\}$ such that $\mathcal{B}_j \cap \mathcal{B}_i \neq \varnothing$ 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, $s_k \in \mathcal{B}_1 \cap \cdots \cap \mathcal{B}_k \cap \mathcal{B}_{k+1}$ and so $\mathcal{B}_j \cap \mathcal{B}_{k+1} \neq \varnothing$ for all $j \leq k$, therefore $s_j, s_{k+1} \in \mathcal{B}_j \cap \mathcal{B}_{k+1}$ and hence $s_1, \ldots, s_{k+1} \in \mathcal{B}_1 \cap \cdots \cap \mathcal{B}_{k+1}$ and (D, u) provides unique minimal support for $\mathcal{B}_{k+1}$ attained at s_{k+1}. So $s_1, \ldots, s_n \in \mathcal{B}_1 \cap \cdots \cap \mathcal{B}_n$ 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 $\bigcap_{i=1}^n \mathcal{B}_i \neq \varnothing$ we have in particular that $\bigcap_{i=1}^n v_i\mathcal{H}[A_i] \neq \varnothing$. By induction and using (3.12) 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 $\mathsf{CE}(H, \mathcal{L}; \mathbb{P}_A)$ 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 $\mathsf{CL}(H[A], \{A_1, \ldots, A_n\})$. Corollary 3.12 now implies that $\mathcal{B} = \bigcup_{i=1}^n \mathcal{B}_i$ is isomorphic with the B-cluster $\mathsf{CL}(H[B], \{B \cap A_1, \ldots, B \cap A_n\})$. $\square$

The case $|A| = 2$, which is not handled in Lemma 4.3, is actually trivial.

Proposition 4.4. *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 $^{\subsetneq}A$-coset extension and the full $^{\subsetneq}A$-coset extension $\mathsf{CE}(H, \mathcal{L}; \mathbb{P}_A)$ has the cluster property.*

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

Combination of this with Lemma 4.3 implies the following result; it encapsulates the first of the two inductive procedures discussed above.

Theorem 4.5 (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 $^{\subsetneq}A$-coset extension. Assume that for all $B \subsetneq A$ and* *every vertex* $v \in \mathcal{L}$ *the full* ${}^{\subsetneq}B$*-coset extension* $\mathsf{CE}(H, v\mathcal{L}[B]; \mathbb{P}_B)$ *embeds into* $v\mathcal{H}[B]$ *and is bridge-free; then the full* ${}^{\subsetneq}A$*-coset extension* $\mathsf{CE}(H, \mathcal{L}; \mathbb{P}_A)$ *has the cluster property.*

Proof. In order to reduce the claim of the theorem to Lemma 4.3, we merely need to argue that the graphs $\mathsf{CE}(H, v\mathcal{L}[B]; \mathbb{P}_B)$ 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 $\mathsf{CE}(H, v\mathcal{L}[B]; \mathbb{P}_B)$ has the cluster property by Proposition 4.4. For $|A| > 3$ we can use the inductive claim for all $|B| < |A|$ (in the rôle of A) to find that $\mathsf{CE}(H, v\mathcal{L}[B]; \mathbb{P}_B)$ 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 4.6. Suppose that a Cayley graph $\mathcal{G}$ covers a complete connected graph $\mathcal{C}$ via a canonical morphism $\varphi\colon \mathcal{G} \twoheadrightarrow \mathcal{C}$ 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 $u \longrightarrow v$ labelled $p \in \widetilde{E}^*$ 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 4.7 (upward induction)**.** *Let* $1 \le k < |E|$ *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* ${}^{\subsetneq}A$*-coset extension (with respect to* H*),*
(2) *the full* ${}^{\subsetneq}A$*-coset extension* $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ *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* $\overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A) \textcircled{\scriptsize v} \mathcal{H}[B]}$ *for* $B \subsetneq A$ *and* v *a vertex of* $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ *(thus, in particular,* $\mathcal{G}$ *covers the graph* $\overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ *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* ${}^{\subsetneq}A$*-coset extension (with respect to* G*),*
(ii) *the full* ${}^{\subsetneq}A$*-coset extension* $\mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ *embeds into* $\mathcal{G}[A]$*,*
(iii) *the embedded full* ${}^{\subsetneq}A$*-coset extension* $\mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ *is bridge-free in* $\mathcal{G}[A]$*.*

Proof. As for (i), that $\mathcal{L}_G$ is admissible for ${}^{\subsetneq}A$-coset extension follows from the fact that $\mathcal{L}_H$ is admissible for ${}^{\subsetneq}A$-coset extension and that the canonical morphism $G \twoheadrightarrow H$ is k-stable. In this case, the canonical morphism $\varphi\colon \mathcal{G} \twoheadrightarrow \mathcal{H}$ is injective on B-components for $B \subsetneq A$ so that condition (3.7) 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 $\psi\colon \mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A) \to \mathcal{G}[A]$ be the canonical graph morphism of Proposition 3.18. We first show that for every $B \subsetneq A$ the restriction of ψ to $\mathsf{CE}(G, \mathcal{L}_G; B)$ 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 $\mathsf{CE}(G, \mathcal{L}_G; B)$ 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 $\overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$; so there is

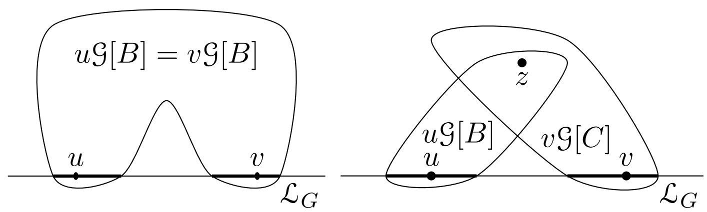

Figure 9

a canonical graph morphism $\varphi\colon \mathcal{G} \twoheadrightarrow \overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ 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 $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$, 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 $u \longrightarrow v'$ 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 $\mathsf{CE}(G, \mathcal{L}_G; B)$ coincide.

So, for the injectivity claim of (ii), it remains to consider the case when vertices of distinct coset extension $\mathsf{CE}(G, \mathcal{L}_G; B)$ and $\mathsf{CE}(G, \mathcal{L}_G; C)$ would violate injectivity. Let $B, C \subsetneq A$, $B \neq C$ and $x \in \mathsf{CE}(G, \mathcal{L}_G; B)$, $y \in \mathsf{CE}(G, \mathcal{L}_G; C)$ be vertices such that $\psi(x) = \psi(y)$. We need to show that $x = y$ in $\mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ (that is, x and y both are in $\mathsf{CE}(G, \mathcal{L}_G; B \cap C)$ 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 $\mathcal{G} \twoheadrightarrow \overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ (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 $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$, 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 $\mathcal{L}_H$). This situation is depicted in Figure 10. So there are paths $u' \xrightarrow{p} c$ in $\mathcal{L}_H \cap u'\mathcal{H}[B]$, $v' \xrightarrow{q} c$ 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,

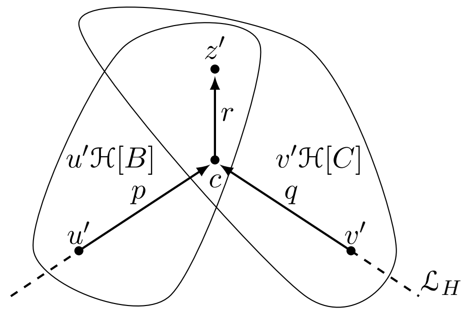

FIGURE 10

pr labels a path $u' \longrightarrow z'$, qr labels a path $v' \longrightarrow z'$. From k-stability of the expansion $G \twoheadrightarrow H$ it follows that the morphism $\mathcal{G} \twoheadrightarrow \overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ 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 $u \longrightarrow z$ labelled pr and one $v \longrightarrow z$ labelled qr. It follows that, in $\mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$,

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

Since $p\colon u' \longrightarrow c$ runs in $\mathcal{L}_H$ and so does $q\colon v' \longrightarrow c$, the path $p\colon u \longrightarrow z \cdot r^{-1}$ runs in $\mathcal{L}_G$, and so does the path $q\colon v \longrightarrow z \cdot r^{-1}$. 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 $\mathsf{CE}(G, \mathcal{L}_G; \{B, C\})$:

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

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

It remains to argue for (iii), that $\mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ 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 $s_1 \in v_1\mathcal{G}[A_1]$, $s_2 \in v_2\mathcal{G}[A_1]$, 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 $\mathcal{L}_G$ and A_i-paths $v_i \xrightarrow{f_i} s_i$ running in $v_i\mathcal{G}[A_i]$. Consider the canonical graph morphism $\varphi\colon \mathcal{G} \twoheadrightarrow \overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$, 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 $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$. Further, the path $s_1 \xrightarrow{p} s_2$ is mapped to the path $s'_1 \xrightarrow{p} s'_2$, which runs in $\overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$.

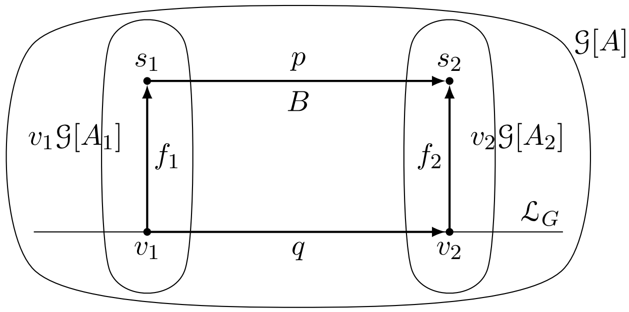

FIGURE 11

It follows that there is a B-path $s'_1 \xrightarrow{p^\circ} s'_2$ which runs in $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ (in fact, p° is the word obtained from p by deletion of the letters which traverse loop edges of $\overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)} \setminus \mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$).

So consider the B-component $\mathcal{B}$ of $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ which contains the two vertices s'_1 and s'_2. The cluster property of $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ 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 $\mathcal{B} = \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_n$ 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 $v'_1\mathcal{H}[C_1] \cap v'_2\mathcal{H}[C_2] \neq \varnothing$ 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

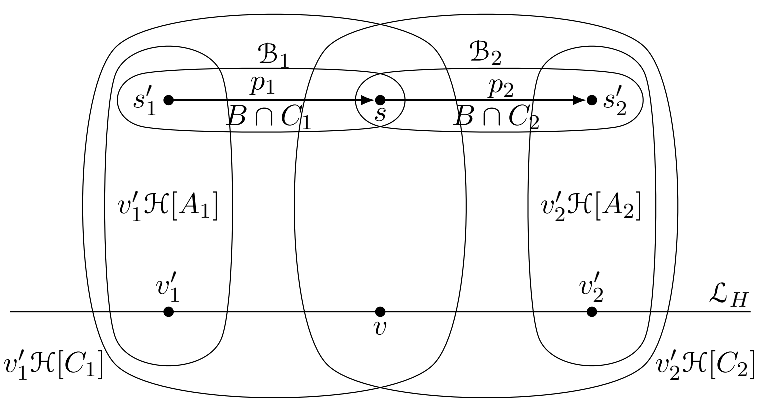

FIGURE 12

$s \in s_1'\mathcal{H}[B \cap C_1] \cap s_2'\mathcal{H}[B \cap C_2]$ 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 $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ at the vertex s and the canonical graph morphism

$$\psi \colon \mathcal{G} \twoheadrightarrow \overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A) ⓢ \mathcal{H}[B]}$$

which maps the covering graph $\mathcal{L}_G$ onto $\mathcal{L}_H$. The graphs $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ and $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)ⓢ\mathcal{H}[B]$ are almost the same except that the cluster $\mathcal{B}$ in the coset extension $\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)$ 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 $[p]_H = [p_1p_2]_H$. Since the expansion $G \twoheadrightarrow H$ is k-stable and $|B| \leq k$, it follows that $[p]_G = [p_1p_2]_G$. 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 $v_1\mathcal{G}[C_1] \cup v_2\mathcal{G}[C_2] \twoheadrightarrow v_1'\mathcal{H}[C_1] \cup v_2'\mathcal{H}[C_2]$ (see Lemma 3.10). It follows that the path $s_1 \xrightarrow{p_1} s \cdot p_1$ runs in $v_1\mathcal{G}[C_1]$ while $s_1 \cdot p_1 \xrightarrow{p_2} s_1 \cdot p_1p_2 = s_2$ runs in $v_2\mathcal{G}[C_2]$. So the path $s_1 \xrightarrow{p_1p_2} s_2$ runs entirely in $v_1\mathcal{G}[C_1] \cup v_2\mathcal{G}[C_2] \subseteq \mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$ and thus provides a B-path between s_1 and s_2 in the coset extension $\mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$. 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 $s_1 \xrightarrow{p_1p_2} s_2$ runs in $v_1\mathcal{G}[C_1]$ which is contained in $\mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$.

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 $\varphi \colon \mathcal{G} \twoheadrightarrow \overline{\mathsf{CE}(H, \mathcal{L}_H; \mathbb{P}_A)}$ induces an isomorphism $\phi \colon s_1\mathcal{G}[B] = s_2\mathcal{G}[B] \twoheadrightarrow v\mathcal{H}[B]$ where $\phi = \varphi \restriction s_1\mathcal{G}[B]$ 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 $\mathsf{CE}(G, \mathcal{L}_G; \mathbb{P}_A)$. □

5. Construction of the group G

The group G announced in Lemma 2.5 will be constructed via a series of expansions

$$G_1 \twoheadleftarrow H_1 \twoheadleftarrow G_2 \twoheadleftarrow \cdots \twoheadleftarrow G_{|E|-1} \twoheadleftarrow H_{|E|-1} \twoheadleftarrow G_{|E|} = G \tag{5.1}$$

where, for every k, the expansions $G_k \twoheadleftarrow H_k$ and $H_k \twoheadleftarrow G_{k+1}$ 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

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

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,

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

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 $\mathcal{E} = (V, \widetilde{E}; \alpha, \omega, {}^{-1})$, here considered as an E-labelled graph where every edge gets its own label. This serves to initialise the series (5.1) with $G_1 := \mathscr{T}(\mathcal{X}_1)$.

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 $H_k \rightsquigarrow G_{k+1}$. The latter step is intended to ensure the following: suppose that p is a word over $k+1$ letters which forms a path $u \longrightarrow v$ in $\mathcal{E}$ and $a \in \mathrm{co}(p)$ for some $a \in E$; if H_k satisfies the relation $p = p_{a\to 1}$, but there is no word q in the letters $B := \mathrm{co}(p) \setminus \{a\}$ (and their inverses) forming a path $u \longrightarrow v$ 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 $p = p_{a\to 1}$ and therefore *every* relation $p = q$ with $q \in \widetilde{B}^*$.

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 $\mathcal{E} = (V, \widetilde{E}, \alpha, \omega, {}^{-1})$ 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 in V that swaps the two vertices αe and ωe. Let $\mathcal{E} = (V, \widetilde{E}; \alpha, \omega, {}^{-1})$ 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 $(V, \widetilde{E}; \alpha, \omega, {}^{-1}, \ell, E)$ where ℓ is the identity function mapping every $e \in \widetilde{E}$, 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

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

We have thus completed every non-loop edge $u \underset{e}{\longrightarrow} v$ to a 2-cycle $u \underset{e}{\overset{e}{\rightleftarrows}} v$.

Let us denote the set of all positive edges so obtained (the original ones and the added ones) by D; then the oriented E-graph $\mathcal{D} = (V, \widetilde{D}; \alpha, \omega, {}^{-1}, \ell, E)$ is weakly complete. Let $\mathcal{X}_1 := \overline{\mathcal{D}}$ be its trivial completion. The transition group $G_1 := \mathscr{T}(\mathcal{X}_1)$ 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 $G_k = \mathscr{T}(\mathcal{X}_k)$. We set

$$\mathcal{Y}_k := \mathcal{X}_k \sqcup \bigsqcup \{\overline{\mathcal{G}_k[A]} \colon A \subseteq E, |A| = k\}. \tag{5.3}$$

Provided that G_k is k-retractable, the transition group $H_k = \mathscr{T}(\mathcal{Y}_k)$ is $(k+1)$-retractable and the expansion $G_k \twoheadleftarrow H_k$ 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 $H_k \twoheadleftarrow G_{k+1}$ 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 $G_{k+1} = \mathscr{T}(\mathcal{X}_{k+1})$. 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 $\textsc{Cond}_k$ for the pair (G_k, H_{k-1}). So $\textsc{Cond}_k$ 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 $\textsc{Cond}_{k+1}$ 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 $\textsc{Cond}_k$, for $k \geq 2$, we collect the following:

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

and, for every $B \subseteq E$ with $|B| \leq k$, for any $\mathcal{G}_k$-cover $\mathcal{C}_{\mathcal{G}_k}$ of any 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 ${}^{\subsetneq}B$-coset extension,
(iii) the full ${}^{\subsetneq}B$-coset extension $\mathsf{CE}(G_k, \mathcal{C}_{\mathcal{G}_k}; \mathbb{P}_B)$ embeds into $\mathcal{G}_k[B]$,
(iv) the embedded full ${}^{\subsetneq}B$-coset extension $\mathsf{CE}(G_k, \mathcal{C}_{\mathcal{G}_k}; \mathbb{P}_B)$ is bridge-free.

By Theorem 3.8, (i) implies that $H_k \twoheadrightarrow G_k$ in particular is k-stable and H_k is $k+1$-retractable by construction, as already mentioned in connection with the definition of H_k. Let $\psi_k \colon \mathcal{G}_k \twoheadrightarrow \mathcal{X}_1$ be some canonical graph morphism, $\chi_k \colon \mathcal{H}_k \twoheadrightarrow \mathcal{G}_k$ 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 $|A| = k+1$ (positive) edges of $\mathcal{E} \subseteq \mathcal{X}_1$ and $\mathcal{A} = \langle A \rangle$ 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 $\varphi_k^{-1}(\mathcal{C})$. We show that $\mathcal{C}_{\mathcal{H}_k}$ is admissible for ${}^{\subsetneq}A$-coset extension (with respect to H_k) and that the full $\subsetneq\!A$-coset extension $\mathsf{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A)$ has the cluster property. From this it will follow that augmented coset extensions of the form $\mathsf{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A) \circledv \mathcal{H}_k[B]$ 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 $\varphi_k(\mathcal{U}) \subseteq \mathcal{C}$ is a B-component of $\mathcal{C}$ and hence is a connected component of $\langle B \rangle \subseteq \mathcal{C}$.

$$\begin{tikzcd}
\mathcal{U} \arrow[d, two heads] & \subseteq & \mathcal{C}_{\mathcal{H}_k} \arrow[d, two heads] & \subseteq & \mathcal{H}_k \arrow[d, two heads, "\varphi_k"'] \\
\varphi_k(\mathcal{U}) & \subseteq & \mathcal{C} & \subseteq & \mathcal{X}_1
\end{tikzcd}$$

By the inductive hypothesis, any $\mathcal{G}_k$-cover $\mathcal{U}'$ of $\varphi_k(\mathcal{U})$ is admissible for $\subsetneq\!B$-coset extension (with respect to G_k) and $\mathsf{CE}(G_k, \mathcal{U}'; \mathbb{P}_B)$ embeds into $\mathcal{G}_k[B]$ and is bridge-free by (ii)–(iv). Since the morphism $\chi_k \colon \mathcal{H}_k \twoheadrightarrow \mathcal{G}_k$ is injective on B-components (that is, injective on B-cosets), it follows that $\mathcal{U}' \cong \mathcal{U}$ and hence also

$$\mathsf{CE}(G_k, \mathcal{U}'; \mathbb{P}_B) \cong \mathsf{CE}(H_k, \mathcal{U}; \mathbb{P}_B). \tag{5.4}$$

Altogether, by (iii) we have

$$\begin{tikzcd}
\mathsf{CE}(H_k, \mathcal{U}; \mathbb{P}_B) \arrow[d, phantom, "\cong" rotate=90] & \mathcal{H}_k[B] \arrow[d, phantom, "\cong" rotate=90] \\
\mathsf{CE}(G_k, \mathcal{U}'; \mathbb{P}_B) \arrow[r, hook] & \mathcal{G}_k[B]
\end{tikzcd}$$

so that $\mathsf{CE}(H_k, \mathcal{U}; \mathbb{P}_B)$ canonically embeds into $\mathcal{H}_k[B]$. It follows that condition (3.7) of Definition 3.16 is fulfilled. Since this is true for every B-component $\mathcal{U}$ for every proper subset B of A this implies that $\mathcal{C}_{\mathcal{H}_k}$ is admissible for $\subsetneq\!A$-coset extension (with respect to H_k). Once more by the inductive hypothesis (iv), every graph in (5.4) is bridge-free. Then, by Theorem 4.5, the full $\subsetneq\!A$-coset extension $\mathsf{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A)$ itself has the cluster property. As already mentioned, this guarantees that the augmented coset extensions $\mathsf{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A) \circledv \mathcal{H}_k[B]$ of Definition 5.3 (2) below are well defined. We therefore 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

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

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

(2) all augmented full $\subsetneq\!A$-coset extensions

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

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 $\mathsf{CE}(H_k, \mathcal{C}_{\mathcal{H}_k}; \mathbb{P}_A)$ and $B \subsetneq A$.

We note that the augmented clusters and augmented coset extensions contain, for $B = \varnothing$, all "plain" clusters and coset extensions. Recall that $\mathcal{H}_k = \mathscr{T}(\mathcal{Y}_k)$ and $G_{k+1} = \mathscr{T}(\mathcal{X}_{k+1}) = \mathscr{T}(\mathcal{Y}_k \sqcup \overline{\mathcal{Z}_k})$; see (5.3) for $\mathcal{Y}_k$.

Proposition 5.4. *The expansion $H_k \twoheadleftarrow G_{k+1}$ 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 $p \in \widetilde{C}^*$ 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 $\mathcal{Y}_k$ or of $\overline{\mathcal{Z}_k}$ 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 $\mathcal{Y}_k$, then we are done since then $[p]_{H_k} \neq 1$ immediately follows from $H_k = \mathscr{T}(\mathcal{Y}_k)$. If $\mathcal{L}$ is a component of $\overline{\mathcal{Z}_k}$, then $\mathcal{L} = \overline{\mathcal{M}}$ where $\mathcal{M}$ is of the form (1) or (2) in Definition 5.3, and the path $p\colon v \longrightarrow v \cdot p$ 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 with one of the following:

(i) a full C-coset $\mathcal{H}_k[C]$, or
(ii) a C-cluster $\mathsf{CL}(H_k[C], \mathbb{P})$ 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 $\mathsf{CL}(H_k[C], \mathbb{P})ⓤ\mathcal{H}_k[D]$ for some set $\mathbb{P}$ of proper subsets of C, some vertex u of $\mathsf{CL}(H_k[C], \mathbb{P})$ 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

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

where the second isomorphism is obvious since $H_k[C] \cong G_k[C]$ by k-stability of $H_k \twoheadrightarrow G_k$ while the third isomorphism follows from Lemma 3.10. In case (iii) we get

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

where t and s are the images of u under the canonical morphisms $H_k \twoheadrightarrow G_k$ and $H_k \twoheadrightarrow H_{k-1}$, 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 $H_k \twoheadrightarrow G_k$ while the third isomorphism follows from Lemma 3.14. Hence, in cases (ii) and (iii), $\overline{v\mathcal{M}[C]}$ 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 4.7 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 $\psi_{k+1} \colon \mathcal{G}_{k+1} \twoheadrightarrow \mathcal{X}_1$ is a canonical graph morphism) is admissible for ${}^{\subsetneq}A$-coset extension, and the full ${}^{\subsetneq}A$-coset extension $\mathsf{CE}(G_{k+1}, \mathcal{C}_{\mathcal{G}_{k+1}}; \mathbb{P}_A)$ 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 ${}^{\subsetneq}A$-coset extension $\mathsf{CE}(G_l, \mathcal{C}_{\mathcal{G}_l}; \mathbb{P}_A)$ embeds into $\mathcal{G}_l[A]$. But the expansion $G_l \twoheadleftarrow G_{k+1}$ is l-stable whence $\mathsf{CE}(G_l, \mathcal{C}_{\mathcal{G}_l}; \mathbb{P}_A) \cong \mathsf{CE}(G_{k+1}, \mathcal{C}_{\mathcal{G}_{k+1}}; \mathbb{P}_A)$ 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 $\textsc{Cond}_{k+1}$:

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

and, for every $A \subseteq E$ with $|A| \leq k+1$, for any $\mathcal{G}_{k+1}$-cover $\mathcal{C}_{\mathcal{G}_{k+1}}$ of every connected component $\mathcal{C}$ of $\mathcal{A} = \langle A \rangle$ in $\mathcal{E} \subseteq \mathcal{X}_1$, the following hold:

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

We check that the base case for this inductive procedure, $\textsc{Cond}_2$ for the pair (G_2, H_1), goes through. The group H_1 is 2-retractable and so is G_2 since $G_2 \twoheadrightarrow H_1$ is 1-stable (cf. Theorem 3.8). By Proposition 4.4, 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 ${}^{\subsetneq}A$-coset extension (with respect to H_1) and $\mathsf{CE}(H_1, \mathcal{C}_{\mathcal{H}_1}; \mathbb{P}_A)$ has the cluster property. Theorem 4.7 then implies that the $\mathcal{G}_2$-cover $\mathcal{C}_{\mathcal{G}_2}$ is admissible for ${}^{\subsetneq}A$-coset extension (with respect to G_2) and that $\mathsf{CE}(G_2, \mathcal{C}_{\mathcal{G}_2}; \mathbb{P}_A)$ 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 $\textsc{Cond}_2$ are satisfied by the pair (G_2, H_1). Altogether the series of expansions

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

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 2.5. 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 $u \longrightarrow v$ in $\mathcal{E}$ there is a G-equivalent word which also forms a path $u \longrightarrow v$ 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 $\mathcal{E} = (V, \widetilde{E}; \alpha, \omega, ^{-1})$. 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 $E \to E$ 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 $p \in \widetilde{E}^*$, we have $p\colon u \longrightarrow v$ if and only if ${}^{\gamma}p\colon {}^{\hat{\gamma}}u \longrightarrow {}^{\hat{\gamma}}v$. Consequently, for every word $p \in \widetilde{E}^*$, 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}. □

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 4.5 and 4.7 proceed by cardinality of subsets of E and treat all subsets of the same size uniformly and in parallel.[3] Any 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

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

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 2.5. Recall that for a word $p \in \widetilde{E}^*$, $\mathrm{co}(p)$ 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). The reader is invited to recall the group G defined in (5.1), the graphs $\mathcal{X}_{k+1} := \mathcal{Y}_k \sqcup \overline{\mathcal{Z}_k}$ (for $\mathcal{Z}_k$ see Definition 5.3) and the coset extensions $\mathsf{CE}(G, \mathcal{K}; \mathbb{P})$ defined in (3.10); the full coset extension $\mathsf{CE}(G, \mathcal{K}; \mathbb{P}_A)$ is defined immediately before Remark 3.17. Also recall that the Cayley graph $\mathcal{G}$ of G covers, in the sense of Definition 3.1, any connected component of any one of the graphs $\mathcal{X}_k$.

Lemma 5.6. *Let $p \in \widetilde{E}^*$ be a word that forms a path $u \longrightarrow v$ in $\mathcal{E}$; let $A = \mathrm{co}(p)$ and suppose that for some letter $a \in A$ and $B = A \setminus \{a\}$ there exists a word $r \in \widetilde{B}^*$ such that $[p]_G = [r]_G$. Then there exists a word $q \in \widetilde{B}^*$ such that $[p]_G = [q]_G$ and, in addition, q forms a path $u \longrightarrow v$ 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 $u \longrightarrow u$ 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 $\varphi_u \colon \mathcal{H}_k \twoheadrightarrow \mathcal{X}_1$ 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 $^{\subsetneq}A$-coset extension $\mathsf{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)$ and note that $\mathsf{CE}(H_k, \mathcal{A}_k; B)$ 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 $\psi \colon \mathcal{G} \twoheadrightarrow \overline{\mathsf{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)}$ (mapping $1 \in \mathcal{G}$ to $1 \in \mathcal{A}_k$) maps $\pi_1^{\mathcal{G}}(p)$ to $\pi_1^{\overline{\mathsf{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)}}(p)$, 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 $r \in \widetilde{B}^*$. 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 $\pi_1^{\overline{\mathsf{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)}}(r)$. But the B-component of 1 in $\overline{\mathsf{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)}$ is the full B-coset $1\mathcal{H}_k[B]$, which is contained in

$\mathsf{CE}(H_k, \mathcal{A}_k; B)$. So the latter graph contains a path labelled r starting at 1, and that path $\pi_1^{\mathsf{CE}(H_k,\mathcal{A}_k;B)}(r)$ 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

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

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 $\pi\colon 1 \longrightarrow v'$ running in $\mathcal{A}_k \cap 1\mathcal{H}_k[B]$; let $q \in \widetilde{B}^*$ be the label of that path. By construction, $[q]_{H_k} = [r]_{H_k}$, hence $[q]_G = [r]_G$ since the expansion $H_k \twoheadleftarrow G$ is k-stable, and therefore also $[q]_G = [p]_G$. Finally, the canonical morphism $\varphi_u\colon \mathcal{H}_k \twoheadrightarrow \mathcal{X}_1$ (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 $u \longrightarrow v$. Altogether, q forms a path $u \longrightarrow v$ 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 $\mathrm{co}(p) = A$ and $|A| = k+1$ such that p forms a path $u \longrightarrow v$ in $\mathcal{E}$, and some letter $a \in A$ does not belong to the $H_k[A]$-content of p then the subgraph $\mathsf{CE}(H_k, \mathcal{A}_k; B)$ of $\overline{\mathsf{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)}$ (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 $r \in \widetilde{B}^*$ (compare Remark 3.19) unless there exists a word $q \in \widetilde{B}^*$ such that $[p]_{H_k} = [q]_{H_k}$ and q forms a path $u \longrightarrow v$ 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 $\mathsf{CE}(H_k, \mathcal{A}_k; B)$ of the mentioned kind (after making them weakly complete by extending edges to 2-cycles whenever needed).[4] 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 $H_k \twoheadleftarrow G_{k+1}$, and it is not clear whether or not k-stability can be achieved by this procedure. Hence, except for the graphs $\mathsf{CE}(H_k, \mathcal{A}_k; B)$, which appear as subgraphs of the full coset extensions $\mathsf{CE}(H_k, \mathcal{A}_k; \mathbb{P}_A)$, 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 $\varphi_u\colon \mathcal{G} \twoheadrightarrow \mathcal{X}_1$ 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 should recall Definition 3.4.

[4]This means that only "basic" B-coset extensions of type $\mathsf{CE}(G, \mathcal{K}; B)$ as in (3.8) would be sufficient for proving Lemma 5.6.

Corollary 5.7. *Let $p \in \widetilde{E}^*$ be a word which forms a path $u \longrightarrow v$ in $\mathcal{E}$; then there exists a word $q \in \widetilde{E}^*$ which uses only letters (i.e. edges) from the content* $\mathrm{C}([p]_G)$ *(and/or their inverses) such that $[p]_G = [q]_G$ and q forms a path $u \longrightarrow v$ in $\mathcal{E}$. If $\mathrm{C}([p]_G) = \varnothing$, then $u = v$ and q is the empty word. If $\mathrm{C}([p]_G) \neq \varnothing$, then the graph $\langle \mathrm{C}([p]_G)\rangle = \langle \mathrm{co}(q)\rangle$ is connected and contains the vertices u and v.*

References

[1] J. Almeida, *Finite Semigroups and Universal Algebra*, World Scientific, Singapore, 1994.

[2] J. Almeida and M. Delgado, Sur certains systèmes d'équations avec contraintes dans un groupe libre, *Portugal. Math.* **56** (1999), no. 4, 409–417.

[3] 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.

[4] N. Alon, Tools from higher algebra. pp. 1749–1783 in R. Graham, M. Grötschel and L. Lovász, eds., *Handbook of Combinatorics*, Vol. II, North-Holland, 1995.

[5] 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.

[6] 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.

[7] K. Auinger and M. B. Szendrei, On F-inverse covers of inverse monoids, *J. Pure Appl. Algebra* **204** (2006), 493–506.

[8] K. Auinger and A. Bors, Formations of finite groups with the M. Hall property, *J. Comb. Algebra* **3** (2019), no. 2, 113–152.

[9] 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.

[10] K. Auinger, J. Bitterlich and M. Otto, Finite approximation of free groups II: The theorems of Ash, Herwig–Lascar and Ribes–Zalesskii — revisited and strengthened. Preprint, in preparation.

[11] 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.

[12] J. Bitterlich, Finite F-inverse covers do exist, preprint (2018), arXiv:1808.05161.

[13] J. Bitterlich, Investigations into the Universal Algebra of Hypergraph Coverings and Applications. Dissertation, Technische Universität Darmstadt, 2019.

[14] T. Coulbois and A. Khelif, Equations in free groups are not finitely approximable, *Proc. Amer. Math. Soc.* **127** (1999), no. 4, 963–965.

[15] G. Z. Elston, Semigroup expansions using the derived category, kernel, and Malcev products, *J. Pure Appl. Algebra* **136** (1999), 231–265.

[16] 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.

[17] 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.

[18] I. Kapovich and A. Myasnikov, Stallings foldings and subgroups of free groups, *J. Algebra* **248** (2002), no. 2, 608–668.

[19] M. V. Lawson, Almost factorisable inverse semigroups, *Glasg. Math. J.* **36** (1994), 97–111.

[20] M. V. Lawson, *Inverse Semigroups. The Theory of Partial Symmetries*, World Scientific, River Edge, NJ, 1998.

[21] 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.

[22] M. Otto, Highly acyclic groups, hypergraph covers, and the guarded fragment, *J. ACM* **59** (2012), no. 1, Art. 5, 40 pp.

[23] M. Otto, Groupoids, hypergraphs, and symmetries in finite models. 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2013), 53–62, IEEE Computer Soc., Los Alamitos, CA, 2013.

[24] M. Otto, Finite groupoids, finite coverings and symmetries in finite structures, preprint (2014), arXiv:1404.4599.

[25] M. Otto, Acyclicity in finite groups and groupoids, preprint (2024), arXiv: 1806.08664v6.

[26] M. Petrich, *Inverse Semigroups*, Wiley, New York, 1984.

[27] L. Ribes and P. Zalesskiĭ, On the profinite topology on a free group, *Bull. London Math. Soc.* **25** (1993), no. 1, 37–43.

[28] J.-P. Serre, *Trees.* Translated from the French original by John Stillwell. Corrected 2nd printing of the 1980 English translation. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2003.

[29] J. R. Stallings, Topology of finite graphs, *Invent. Math.* **71** (1983), no. 3, 551–565.

[30] N. Szakács, On the graph condition regarding the F-inverse cover problem, *Semigroup Forum* **92** (2016) no. 3, 551–558.

[31] N. Szakács and M. B. Szendrei, On F-inverse covers of finite-above inverse monoids, *J. Algebra* **452** (2016), 42–65.

[32] M. B. Szendrei, Some open problems in the structure theory of regular semigroups, *Semigroup Forum* **64** (2002), no. 2, 213–223.

[33] 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

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

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

Department of Mathematics, Technische Universität Darmstadt, Schlossgartenstrasse 7, D-64289 Darmstadt, Germany

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