Acyclicity in Finite Groups and Groupoids

Martin Otto*
Department of Mathematics
Technische Universität Darmstadt

January 2020

Abstract

We expound a concise construction of finite groups and groupoids whose Cayley graphs satisfy graded acyclicity requirements. Our acyclicity criteria concern cyclic patterns formed by coset-like configurations w.r.t. subsets of the generator set rather than just by individual generators. The proposed constructions correspondingly yield finite groups and groupoids whose Cayley graphs satisfy much stronger acyclicity conditions than large girth. We thus obtain generic and canonical constructions of highly homogeneous graph structures with strong acyclicity properties, which support known applications in finite graph and hypergraph coverings that locally unfold cyclic configurations. with involutive generators, with the additional benefit of a more uniform approach across these settings.

*Research partially supported by DFG grant OT 147/6-1: *Constructions and Analysis in Hypergraphs of Controlled Acyclicity.*

Contents

1 Introduction

The intimate connection between finite groups and graph-like structures is a long-standing theme that illustrates core concepts at the interface of algebra and discrete mathematics. Groups arise as automorphism groups of structures, and Frucht's theorem [10] says that every finite group arises as an automorphism group of a finite graph; in particular, the given finite group – an abstract group – is realised as a permutation group, and thus as a subgroup of the full symmetric group of some finite set, and in fact even as the full group of all symmetries of a specifically designed discrete structure of a very simple format.

At a very basic level, permutation group actions can be defined through generators whose operation can be traced in graph-like structures, which in turn determine the abstract group structure [6, 7]. The key notion in this correspondence is the representation of the algebraic structure of the given group in its Cayley graph: an edge-coloured directed graph that represents the internal group action of a chosen set of generators for the group.

Interesting finite groups can be obtained as permutation group actions induced by graph-like extra structure on a finite set, also in other ways than just as a group of symmetries. Specific graph structures and carefully designed permutation group actions can thus give rise to finite groups with desirable algebraic or combinatorial properties suggested by various applications. A very nice example of this technique is a construction, due to Biggs [4] and outlined in [1], of finite groups over a given set of generators that avoid short cycles, i.e. in which non-trivial products of a small number of generators cannot evaluate to the neutral element. In terms of the Cayley graph of the resulting group one obtains finite graphs of large girth that are not only regular but (like any Cayley graph) highly symmetric in the stronger sense of possessing a transitive automorphism group.

Acyclicity criteria for groups matter in many natural applications. The free group over a given set of generators, which can be seen as the unique fully acyclic group structure over the given generators, arises naturally in connection with universal coverings in the classical topological context as well as in the context of discrete structures, e.g. with tree unfoldings of transition systems. The relevant coverings can be described as products with (the Cayley graphs of) free groups. Of course free groups, and fully acyclic coverings in non-trivial settings, are necessarily infinite. Where finiteness matters and needs to be preserved, e.g. in finite coverings, full acyclicity is typically unavailable. Here graded degrees of acyclicity, like lower bounds on the girth of the Cayley graph, are best possible and often can replace full acyclicity, especially for local structural analysis – just as a graph of large girth is locally tree-like. Previous work, which arose from applications in logic and the model theory of finite structures, has led to the introduction of similar but much stronger measures of graded acyclicity in Cayley graphs of finite groups. These notions of acyclicity arise naturally in connection with covering constructions for finite graphs and hypergraphs. Instead of controlling just the length of shortest generator cycles, similar control is achieved over the length of shortest cycles formed by cosets w.r.t. generated subgroups. This generalisation involves a passage from cycles at the level of individual generators to cycles formed by cosets, which a priori are not even bounded in size. In other words, this is a shift in focus from first-order objects (generators) to second-order objects (cosets) in the desired groups. Corresponding constructions, which are inspired by Biggs' technique but adapt the basic idea to the more complex technical setting, were first developed for groups in [12] and [13] for specific applications of finite graph coverings. Generalisations of these techniques to the setting of groupoids offer a more direct route to hypergraph coverings (here necessarily branched, in a discrete analogue of classical terminology from [9]). A main challenge and goal in these settings lies in the construction of corresponding coverings that are generic and natural in the sense that they do not break any symmetries of the underlying structure. This is essential for far-reaching applications, e.g. towards extension problems for local symmetries [14, 15].

The goal here is a concise and generic combinatorial construction of groups and groupoids with strong acyclicity properties that control coset cycles rather than just generator cycles. The present exposition not only serves to correct a serious mistake in the construction of the relevant groupoids that was sketched in [14, 15][1] but also to unify the treatment of finite groups and groupoids with the desired acyclicity properties. One main technical point in the generalisation from groups to groupoids has to do with the difficulty to overcome the restriction to involutive generators from [4, 12], which seems inadequate in a groupoidal setting. In the current, more comprehensive and more systematic extension of the original idea we propose a construction of highly acyclic finite groups with sets of involutive generators that yields stronger results for these groups – or stronger notions of acyclicity based on more general patterns than mere coset cycles. This allows us to present a self-contained account in which the construction of highly acyclic finite groupoids can be reduced to the new, enriched construction for groups with involutive generators. This yields a unified construction which offers a transparent view of the commonality between the two, seemingly so very different settings, which may support further insights and applications. Concerning known applications we discuss more general and more direct constructions of finite graph and hypergraph coverings in Propositions 10.1 and 10.2.

Terminology and notation

Graphs and relational structures. In this paper we consider various kinds of graphs, some undirected, some directed, often also allowing loops (reflexive edges), and in Section 8 also multi-graphs that may have more than one edge linking the same two vertices. Notation should be standard, with small adaptations to the specific formats that will be explicitly stated where they occur. We mostly use a relational format for the specification of a graph, with a binary edge relation, or with a separate edge relation for each colour to encode

[1] Cf. acknowledgements at the end of this paper on this somewhat frayed history.

edge-coloured graphs. In some instances, and especially in Sections 8 and 9, it is natural to treat graphs and especially multi-graphs as two-sorted structures with a set of edges and a set of vertices linked by incidence maps that specify source and target vertices of each edge. For *subgraphs* we explicitly distinguish between *induced subgraphs* (whose edge relation is the restriction of the given edge relation to the restricted set of vertices) and *weak subgraphs* (whose edge relation may be a proper subset of the given edge relation even in restriction to the smaller vertex set). Also more generally for relational structures we use $\subseteq_w$ for the *weak substructure* relationship, $\subseteq$ for the induced *substructure* relationship. By a *component* of a graph structure we mean an induced substructure that is closed w.r.t. the edge relation; a *connected component* is a minimal component. The term *reduct* refers to a restriction in the number of edge relations, or edge colours, which corresponds to the deletion of all edges of the colours to be eliminated.[2]

Algebraic structures. For structures like groups, semigroups, monoids or groupoids we adopt multiplicative notation and would typically write, for instance, $g \cdot h$ or just gh for the result of the composition of group elements g and h w.r.t. the group operation, 1 for the neutral element and g^{-1} for the inverse of g. When dealing with subgroups of the symmetric group of some set X, we sometimes make the group operation explicit as in $h \circ g$ for the composition of g with h, which maps $x \in X$ to $h(g(x))$, and would in our standard notation be rendered as $g \cdot h$ or gh (!) since we think of permutations as operating from the right.

Among standard terminology from other fields of mathematics we use some basic terms from formal language theory, especially to deal with *words* over a finite alphabet E of letters; the set of all E-words is the set of all finite (but possibly empty) strings or tuples of letters from E, denoted $E^* = \bigcup_{n \in \mathbb{N}} E^n$. As is common in formal language theory, we write a typical word of length $n \in \mathbb{N}$ as $w = e_1 e_2 \cdots e_n \in E^n$ (rather than e.g., in tuple notation, as $(e_1, e_2, \ldots, e_n)$), denoting its length as $n = |w|$. We also write, e.g. just $w_1 w_2$ for the *concatenation* of the words $w_1, w_2 \in E^*$ (which is often denoted as $w_1 \cdot w_2$ with explicit notation for the concatenation operation as a monoidal semigroup operation). The *empty word* $\lambda \in E^*$, which is the unique E-word of length 0, is the neutral element in the monoid E^*. Depending on the rôle of the letters $e \in E$, we may use E-words to specify different objects of interest: thinking of E as a set of generators of some group, an E-word is a generator word which can be read as a group product specifying a group element; thinking of E as a set of colours in an edge-coloured graph, an E-word is a colour sequence and can specify the class of walks that realise that colour sequence. In some cases we also invoke a notion of *reduced words*, which are typically obtained by some cancellation operation. Especially if E is a set of generators of a group that is closed under inverses

[2]Depending on context the corresponding edge relations can be thought of as erased (which produces a structure over a smaller signature), or just as emptied (which produces a weak substructure).

we may (inductively) cancel factors ee^{-1} in order to associate with every E-word a unique reduced E-word that denotes the same group element. In such contexts we often let E^* stand for the set of reduced words, endowed with the concatenation operation that implicitly post-processes plain concatenation by the necessary cancellation steps. More formally one could explicitly distinguish between E^* and its quotient $E^*/\sim$, but we suppress this as an unnecessary distraction in our considerations.

2 General patterns

2.1 Cayley & Biggs: the basic construction

The fundamental idea to associate groups with permutation group actions and graphs can be attributed to Arthur Cayley [6, 7]. The Cayley graph of an abstract group, w.r.t. to a chosen set of generators, encodes the algebraic structural information about the algebraic group, and also represents the given group as a subgroup of the full symmetric group, and more specifically as the automorphism group, of the Cayley graph. The natural passage between combinatorial properties of graph-like structures and group-like structures offers interesting avenues for the construction of group-like and graph-like structures. A classical example is the use of Cayley graphs in Frucht's construction of (finite) graphs that realise a given abstract (finite) group as their automorphism group [10]. In particular, Cayley graphs are, by construction, not just regular but homogeneous in the sense of having a transitive automorphism group. So on one hand, Cayley graphs provide examples of graph structures with a particularly high degree of internal symmetry. On the other hand, permutation group actions on suitably designed graph structures generate groups that can display specific combinatorial properties w.r.t. to a chosen set of generators – and these groups in turn generate Cayley graphs that reflect those group properties. It is one characteristic feature of the inductive constructions to be expounded here that they are based on a feedback loop built on this interplay.

The idea to extract groups with certain acyclicity properties from permutation group actions on suitably prepared graph structures is best illustrated by the basic example of a construction of regular graphs of high girth due to Biggs [4] and outlined in [1].

Let E be a finite set of letters, $|E| = d \geqslant 2$, to be used to label involutive generators of a group to be constructed. With E and a parameter $n \geqslant 1$ in $\mathbb{N}$ associate a tree $\mathbb{T}(E, n)$ and a group $\mathbb{G}(E, n)$ as follows. Let $\mathbb{T}(E, n)$ be a d-branching, regularly E-coloured, finite undirected tree of depth n, as represented by the set of all *reduced words* $w \in E^{\leqslant n} \subseteq E^*$, i.e. strings $w = e_1 \cdots e_m$ of length $|w| = m$, $0 \leqslant m \leqslant n$, with $e_i \in E$ for $1 \leqslant i \leqslant m$ and $e_{i+1} \neq e_i$ for $1 \leqslant i < m$. We regard the empty word $\lambda \in E^*$ as the root of $\mathbb{T}(E, n)$. More formally, we let

$$\mathbb{T}(E, n) = (V, (R_e)_{e \in E})$$

be the tree structure with vertex set

$$V := \{ w \in E^* \colon |w| \leqslant n, w \text{ reduced } \}$$

and undirected edge relation $R = \dot{\bigcup}_{e \in E} R_e$, E-coloured by its partition into the

$$R_e := \{ (w, we), (we, w) \colon w, we \in V \}$$

for $e \in E$. By construction, each vertex $w \in V$ with $|w| < n$ is an interior vertex of $\mathbb{T}(E, n)$ of degree $d = |E|$, with precisely one R_e-neighbour for each $e \in E$; the remaining vertices, viz. those $w \in V$ with $|w| = n$, are leaves of $\mathbb{T}(E, n)$, each with an R_e-neighbour for a unique $e \in E$ (the last letter of w). Note that each R_e is a partial matching over V, and that R_e and $R_{e'}$ are disjoint for $e \neq e'$. With $e \in E$ we associate the permutation $\pi_e \in \mathrm{Sym}(V)$ that swaps any pair of vertices that are incident with a common e-coloured edge. This is the involutive permutation of V whose graph is the matching R_e augmented by loops in vertices not incident with an e-coloured edge. The target of the construction is the group $\mathbb{G}(E, n)$, which is the subgroup of $\mathrm{Sym}(V)$ generated by these involutions:

$$\mathbb{G} = \mathbb{G}(E, n) := \langle \pi_e \colon e \in E \rangle \subseteq \mathrm{Sym}(V).$$

For the group operation we use the convention that the action by the generators is regarded as a right action via composition, i.e. with

$$\rho\pi_e = \pi_e \circ \rho \colon \begin{array}{rcl} V & \longrightarrow & V \\ w & \longmapsto & \pi_e(\rho(w)). \end{array}$$

Its Cayley graph w.r.t. the generators $(\pi_e)_{e \in E}$ is an edge-coloured graph $\mathbb{CG}$, with the set of group elements $\rho \in G$ as its vertex set, and with a family of edge relations

$$R_e^G := \{ (\rho, \rho\pi_e) \colon \rho \in G, e \in E \} \subseteq G \times G,$$

one for each $e \in E$. Here these edge relations are symmetric due to the involutive nature of the π_e in $\mathrm{Sym}(V)$, and they are irreflexive and pairwise disjoint since $\mathrm{id}_V \neq \pi_e \neq \pi_{e'}$ for $e \neq e'$, as can be seen most easily by their action as permutations on $\lambda \in V$. So this Cayley graph is a d-regular finite graph, whose automorphism group acts transitively on the set of vertices. For the last claim consider the left action of the group on itself:

$$h \colon \begin{array}{rcl} G & \longrightarrow & G \\ g & \longmapsto & hg, \end{array}$$

which clearly induces an automorphism of the Cayley graph (albeit not of the group, which is rigid once we label the generators). That the girth of the Cayley graph of $\mathbb{G}$ is at least $4n + 2$ can be seen as follows. A reduced word $w \in E^k$ of length $k \geqslant 1$ can be written as $w = e_1 u$. Let $v \in E^n$ be a leaf of $\mathbb{T}(E, n)$ whose reversal v^{-1} agrees with u (up to $|u|$). Applying the corresponding permutation

$\pi_w = \pi_u \circ \pi e_1$ to v, we see that the action of the permutations prescribed by the first (up to) $n+1$ letters of w takes that leaf step by step towards the root λ, the next n letters (if present) will take it step by step towards a different leaf, where the very next letter (if present) can have no effect so that it would take at least the action of another $2n$ letters after that to bring this vertex back to where we started. In other words, no reduced word of fewer than $n+1+n+1+2n = 4n+2$ letters can label a generator sequence that represents the neutral element of the group, which is the identity in $\mathrm{Sym}(V)$.

Considering what is essential for the passage from a graph like $\mathbb{T}(E,n)$ to a group like $\mathbb{G}(E,n)$, the only obvious necessity is that each of the edge colours induces a partial matching of the underlying vertex set in order to have well-defined involutions π_e. Tree-likeness, by contrast, is of no special importance, not even for the bound on the girth of the resulting group or Cayley graph. If $\mathbb{T}(E,n)$ were replaced, for instance, by the disjoint union of all E-coloured line graphs corresponding to reduced words $w \in E^{2n}$, the above girth bound of $4n+2$ persists with essentially the same argument. In the following paragraph we extract the basic format for the generation of groups with involutive generators from edge-coloured undirected graphs.

2.2 E-graphs and E-groups

In the following it is convenient to allow loops in the symmetric edge relation of an undirected graph (V,R), and to let a loop at vertex v contribute value 1 to the degree of that vertex. A *partial matching* is here cast as a symmetric edge relation whose degree is bounded by 1 at every vertex, and may thus be thought of as the graph of a partial bijection that is involutive (its own inverse); this involution has precisely those vertices as fixed points at which the edge relation has loops, and its domain $\mathrm{dom}(R)$ and range $\mathrm{rng}(R)$ consists of the set of the vertices of degree 1. A *full matching* is a symmetric edge relation R on V such that every vertex $v \in V$ has a unique R-neighbour, which in the case of a loop may be v itself; it therefore corresponds to the graph of an involutive permutation of the vertex set V.

Definition 2.1. [E-graph]
For a set E, an E-*graph* is an undirected edge-coloured graph $\mathbb{H} = (V, (R_e)_{e \in E})$ whose undirected edges are E-coloured in such a way that each R_e is a partial matching over the vertex set V. The E-graph $\mathbb{H} = (V, (R_e)_{e \in E})$ is *strict* if there are *no loops* (each R_e is irreflexive) and *no multiple edges* ($R_e \cap R_{e'} = \emptyset$ for $e \neq e'$). The E-graph $\mathbb{H} = (V, (R_e)_{e \in E})$ is *complete* if each R_e is a full matching. The *trivial completion* of an E-graph $\mathbb{H} = (V, (R_e)_{e \in E})$ is the complete E-graph $\bar{\mathbb{H}} = (V, (\bar{R}_e)_{e \in E})$ obtained by putting $\bar{R}_e := R_e \cup \{(v,v) : v \in V \setminus \mathrm{dom}(R_e)\}$.

We think of R_e-edges as edges of colour e or as edges labelled with e. In this sense an E-graph is a special kind of E-coloured graph whose overall edge relation would be $\bigcup_{e \in E} R_e$.

For groups $\mathbb{G} = (G, \cdot, 1)$ (in multiplicative notation), an element $g \in G$ is an *involution* if $g = g^{-1}$. A subset $E \subseteq G \setminus \{1\}$ is a *set of generators* for $\mathbb{G}$ if every group element $g \in G$ can be written as a product of elements from E and their inverses.

Definition 2.2. [E-group]
For a set E, an E-*group* is any group $\mathbb{G} = (G, \cdot, 1)$ that has $E \subseteq G$ as a set of non-trivial involutive generators.[3]

If $\mathbb{G}$ is an E-group, we write $[w]_{\mathbb{G}} \in G$ for the group element that is the group product of the generator sequence $w \in E^*$, so that

$$\begin{array}{rcl} [\;]_{\mathbb{G}} \colon E^* & \longrightarrow & \mathbb{G} \\ w = e_1 \cdots e_n & \longmapsto & [w]_{\mathbb{G}} := \prod_{i=1}^n e_i = e_1 \cdots e_n \end{array}$$

is a surjective homomorphism from the free monoid structure of E^*, with concatenation and neutral element $\lambda \in E^*$, onto the group $\mathbb{G}$.

Observation 2.3. *The quotient of the free group generated by E w.r.t. to the equivalence relation induced by the identities $e = e^{-1}$ for $e \in E$ (as represented by reduced words in E^*) can be regarded as the* free E-group*. All other* E-*groups are homomorphic images of this free* E-*group.*

Definition 2.4. [sym($\mathbb{H}$)]
For an E-graph $\mathbb{H} = (V, (R_e)_{e\in E})$ we let sym($\mathbb{H}$) be the subgroup of sym(V) that is generated by the involutive permutations $\pi_e \colon V \to V$ induced by the full matchings of its trivial completion $\bar{\mathbb{H}} = (V, (\bar{R}_e)_{e\in E})$.

Provided the $(\pi_e)_{e\in E}$ are pairwise distinct and distinct from id_V, we regard sym($\mathbb{H}$) as an E-group where we identify $e \in E$ with the generator π_e, for $e \in E$. We shall always tacitly assume this whenever we use sym($\mathbb{H}$). A simple manner to force the necessary distinctions for the π_e is to attach to $\mathbb{H}$, as a disjoint component, a copy of the hypercube 2^E (and this modification will nowhere interfere with other concerns of our constructions).

The Biggs group $\mathbb{G}(E, n)$ as discussed above is sym($\mathbb{T}(E, n)$).

Recall that we let permutations act from the right. In terms of the group product in sym($\mathbb{H}$) this makes $\pi_e\pi_{e'} = \pi_{e'} \circ \pi_e$. Extending this to arbitrary words $w = e_1 \cdots e_n \in E^*$ over E according to

$$\begin{array}{rcl} [\;]_{\mathbb{H}} \colon E^* & \longrightarrow & \mathrm{sym}(\mathbb{H}) \\ w & \longmapsto & [w]_{\mathbb{H}} := \pi_w := \prod_{i=1}^n \pi_{e_i} = \pi_{e_n} \circ \cdots \circ \pi_{e_1}, \end{array}$$

yields a surjective homomorphism from the free monoid structure of E^*, with concatenation and neutral element $\lambda \in E^*$, onto the group structure of sym($\mathbb{H}$) with composition and neutral element $\pi_\lambda = \mathrm{id}_V$. Factorisation w.r.t. the identities $e = e^{-1}$ turns this into a surjective group homomorphism from the free E-group onto sym($\mathbb{H}$).

[3] Clearly the elements $e \in E \subseteq G$ are pairwise distinct as elements of $\mathbb{G}$, and non-triviality means that $e \neq 1$.

Definition 2.5. [Cayley graph]
For an abstract group $\mathbb{G} = (G, \cdot, 1)$ and any set $E \subseteq G$ of generators, the *Cayley graph* of G w.r.t. E is the directed edge-coloured graph $\mathbb{CG} := \mathrm{Cayley}(\mathbb{G}, E) = (G, (R_e)_{e\in E})$ with vertex set G and edge sets

$$R_e := \{(g, ge) \colon g \in G\}$$

of colour e, for all $e \in E$.

The Cayley graph $\mathbb{CG}$ is undirected precisely if the generator set E consists of involutions of $\mathbb{G}$. In general the R_e will not be symmetric, but each R_e will always be the graph of a global permutation π_e of the vertex set G, viz. of right multiplication with $e \in G$, $\pi_e \colon g \mapsto ge$. It is easy to check that, as an abstract group with generators $e \in E$, $\mathbb{G}$ is isomorphic to the subgroup of the full symmetric group $\mathrm{Sym}(G)$ over the vertex set G generated by these permutations π_e. In particular, in the case of a group $\mathbb{G} = (G, \cdot, 1)$ that admits a set of involutive generators $E \subseteq G \setminus \{1\}$, the associated Cayley graph $\mathbb{CG} = \mathrm{Cayley}(\mathbb{G}, E)$ is a complete and strict E-graph in the sense of Definition 2.1, and

$$\mathbb{G} = (G, \cdot, 1) \;\simeq\; \mathrm{sym}(\mathbb{CG}).$$

In the following it will be convenient, and without risk of confusion, to identify the generators $e \in E \subseteq G$ of a group $\mathbb{G}$ with the maps $\pi_e \colon g \mapsto ge$ in $\mathbb{G}$ or in its Cayley graph $\mathbb{CG}$. We similarly identify the family of generators $(\pi_e)_{e\in E}$ of $\mathrm{sym}(\mathbb{H})$ with a subset $E \subseteq \mathrm{sym}(\mathbb{H})$ whenever $\mathrm{sym}(\mathbb{H})$ is an E-group, by writing just e instead of π_e in this context.

Definition 2.6. [generated subgroup]
For a subset $\alpha \subseteq E$ of the set of involutive generators E of an E-group $\mathbb{G}$ we let $\mathbb{G}[\alpha]$ stand for the subgroup generated by α, regarded as an α-group whose universe is

$$G[\alpha] := \{[w]_{\mathbb{G}} \colon w \in \alpha^*\} \subseteq G.$$

The Cayley graph $\mathbb{CG}[\alpha]$ of $\mathbb{G}[\alpha]$, correspondingly, is regarded as an α-graph, which is a weak subgraph $\mathbb{CG}[\alpha] \subseteq_w \mathbb{CG}$ of the Cayley graph of $\mathbb{G}$.[4]

Definition 2.7. [α-walk and α-component]
For a subset $\alpha \subseteq E$ and an E-graph $\mathbb{H} = (V, (R_e)_{e\in E})$, an *$\alpha$-walk* of length n from v to v' is a sequence $v_0, e_1, v_1, e_2, \ldots, e_n, v_n$ of vertices and edge labels where $v_i \in V$, $v = v_0$, $v_n = v'$, $e_i \in \alpha$ such that $(v_i, v_{i+1}) \in R_{e_{i+1}}$ for $i < n$.
The *α-connected component*, or just α-component, of $v \in V$ consists of those vertices v' that are linked to v by α-walks. We write $\alpha[v] \subseteq V$ for this set of vertices and $\mathbb{H}[\alpha; v]$ for the weak subgraph $\mathbb{H}[\alpha; v] \subseteq_w \mathbb{H}$ obtained, as an α-graph, as a reduct of the induced subgraph $\mathbb{H}\restriction\alpha[v]$.

Note that the Cayley graph $\mathbb{CG}[\alpha] \subseteq_w \mathbb{CG}$ of $\mathbb{G}[\alpha]$ also arises as the α-component of $1 \in G$ in the Cayley graph $\mathbb{CG}$. It is also useful to note that, if $v = v_0, e_1, v_1, e_2, \ldots, e_n, v_n = v'$ is an α-walk from v to v' in the E-graph $\mathbb{H}$ such that $w = e_1 \cdots e_n$ traces the edge labels along this walk, then $v' = \pi_w(v) = [w]_{\mathbb{H}}(v)$ w.r.t. the permutation group action of $\mathrm{sym}(\mathbb{H})$ on $\mathbb{H}$.

[4]More specifically it is the $(R_e)_{e\in\alpha}$-reduct of the induced subgraph $\mathbb{CG}\restriction G[\alpha]$ on $G[\alpha] \subseteq G$.

2.3 Compatibility and homomorphisms

The notion of a homomorphism between E-groups is the natural one. It requires compatibility with the group product *and* with the identification of the generators. We write $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$ or $\mathbb{G} \preccurlyeq \hat{\mathbb{G}}$ to indicate that there is a homomorphism from $\hat{\mathbb{G}}$ to $\mathbb{G}$. If there is any homomorphism $h\colon \hat{\mathbb{G}} \to \mathbb{G}$ between E-groups $\hat{\mathbb{G}}$ and $\mathbb{G}$ then it must be

$$\begin{array}{rccc} h\colon & \hat{\mathbb{G}} & \longrightarrow & \mathbb{G} \\ & [w]_{\hat{\mathbb{G}}} & \longmapsto & [w]_{\mathbb{G}} \end{array}$$

for all $w \in E^*$. So what matters is well-definedness of this mapping, which is expressible as the condition that $[w]_{\mathbb{G}} = [u]_{\mathbb{G}}$ whenever $[w]_{\hat{\mathbb{G}}} = [u]_{\hat{\mathbb{G}}}$, or just that $[w]_{\mathbb{G}} = 1$ in $\mathbb{G}$ whenever $[w]_{\hat{\mathbb{G}}} = 1$ in $\hat{\mathbb{G}}$.

Definition 2.8. [compatibility]
For an E-graph $\mathbb{H}$ and E-group $\mathbb{G}$ we say that $\mathbb{G}$ is *compatible with* $\mathbb{H}$ if there is a homomorphism of E-groups from $\mathbb{G}$ to $\mathrm{sym}(\mathbb{H})$, i.e. if $\mathrm{sym}(\mathbb{H}) \preccurlyeq \mathbb{G}$.

In straightforward extension of this concept, a family of E-groups is compatible with $\mathbb{H}$ if each member is; this will be of interest especially when certain families of (small) generated subgroups of $\mathbb{G}$, rather than $\mathbb{G}$ itself, are compatible with some E-graph.

Recall that $[w]_{\mathbb{H}} = \pi_w = \prod_{i=1}^n \pi_{e_i} \in \mathrm{sym}(\mathbb{H})$ for $w = e_1 \cdots e_n \in E^*$. Compatibility of $\mathbb{G}$ with $\mathbb{H}$ precisely requires that the mapping

$$[w]_{\mathbb{G}} \longmapsto [w]_{\mathbb{H}} \in \mathrm{sym}(\mathbb{H}).$$

is well-defined, i.e. that $[w]_{\mathbb{G}} = 1$ in $\mathbb{G}$ implies $[w]_{\mathbb{H}} = 1$ in $\mathrm{sym}(\mathbb{H})$, for all $w \in E^*$.

Note that trivially $\mathrm{sym}(\mathbb{H})$ is compatible with $\mathbb{H}$ and with every connected component of $\mathbb{H}$. We collect some further simple but useful facts.

Observation 2.9. *(i)* $\mathbb{G}[\alpha]$ *is compatible with* $\mathbb{C}\mathbb{G}[\alpha]$ *for* $\alpha \subseteq E$.

(ii) $\mathbb{G}$ *is compatible with the disjoint union* $\oplus_i \mathbb{H}_i$ *of* E*-graphs* $\mathbb{H}_i$ *if, and only if, it is compatible with each component* $\mathbb{H}_i$.

(iii) $\mathbb{G}[\alpha]$*, for* $\alpha \subseteq E$*, is compatible with* $\mathbb{H}$ *if, and only if, it is compatible with every* α*-connected component of* $\mathbb{H}$*, if, and only if, it is compatible with the* α*-reduct of* $\mathbb{H}$.

It also follows that $\mathbb{G} \preccurlyeq \hat{\mathbb{G}}$ if, and only if, $\hat{\mathbb{G}}$ is compatible with $\mathbb{C}\mathbb{G}$. A version of this observation for generated subgroups will be crucial in the construction of suitable E-groups with specific acyclicity properties.

Lemma 2.10. *Let* $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$ *be* E*-groups,* $\hat{\mathbb{G}} = \mathrm{sym}(\mathbb{H})$ *for an* E*-graph* $\mathbb{H}$*. In this situation, the subgroups* $\mathbb{G}[\alpha]$ *and* $\hat{\mathbb{G}}[\alpha]$ *generated by* $\alpha \subseteq E$ *are isomorphic as* α*-groups,* $\hat{\mathbb{G}}[\alpha] \simeq \mathbb{G}[\alpha]$*, if the homomorphism from* $\hat{\mathbb{G}}$ *to* $\mathbb{G}$ *is injective in restriction to* $\hat{\mathbb{G}}[\alpha]$*, which is the case if, and only if,* $\mathbb{G}[\alpha]$ *is compatible with every* α*-connected component of* $\mathbb{H}$ *and hence with* $\mathbb{H}$.

Proof. For the last claim, assuming that $\mathbb{G} \preccurlyeq \hat{\mathbb{G}}$, we need to show that conversely $\hat{\mathbb{G}}[\alpha] \preccurlyeq \mathbb{G}[\alpha]$. Note that $\hat{\mathbb{G}}[\alpha] = \mathrm{sym}(\mathbb{H}\restriction\alpha)$, where $\mathbb{H}\restriction\alpha$ stands for the α-reduct of $\mathbb{H}$, which is an α-graph. If $\mathbb{G}[\alpha]$ is compatible with every connected component of $\mathbb{H}\restriction\alpha$, then $[w]_{\mathbb{G}} = [u]_{\mathbb{G}}$ for $w, u \in \alpha^*$ implies that $\pi_w = \pi_u$ in $\mathrm{sym}(\mathbb{H}\restriction\alpha)$ and therefore also in $\mathrm{sym}(\mathbb{H})$, i.e. in $\hat{\mathbb{G}}$. □

For $k \in \mathbb{N}$ we let

$$\Gamma_k := \{\alpha \subseteq E \colon |\alpha| < k\}$$

denote the set of generator subsets of size less than k. We abbreviate corresponding families of generated subgroups and their Cayley graphs from a given E-group $\mathbb{G}$ as

$$\Gamma_k(\mathbb{G}) := (\mathbb{G}[\alpha'] \colon \alpha' \in \Gamma_k)$$
$$\Gamma_k(\mathbb{C}\mathbb{G}) := (\mathbb{C}\mathbb{G}[\alpha'] \colon \alpha' \in \Gamma_k)$$

Lemma 2.11. *Let $\hat{\mathbb{G}} = \mathrm{sym}(\mathbb{H})$ for an E-graph $\mathbb{H} = \mathbb{H}(\mathbb{G})$ that is derived from the E-graph $\mathbb{G}$ in an isomorphism-respecting manner.*

(a) If $\mathbb{H}(\mathbb{G})$ comprises a copy of $\mathbb{C}\mathbb{G}$ as a component, then $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$.

(b) If $\mathbb{H}(\mathbb{G})$ as in (a) is such that $\Gamma_k(\mathbb{G})$, i.e. every $\mathbb{G}[\alpha]$ for $\alpha \in \Gamma_k$, is compatible with $\mathbb{H}(\mathbb{G})$, then $\hat{\mathbb{G}}[\alpha] \simeq \mathbb{G}[\alpha]$ for $\alpha \in \Gamma_k$.

(c) If $\mathbb{H}(\mathbb{G})$ as in (b) comprises, as a component, an E-graph $\mathbb{H}_k(\mathbb{G})$ whose isomorphism type is determined by the family of the subgroups $\Gamma_k(\mathbb{G}) = (\mathbb{G}[\alpha] \colon \alpha \in \Gamma_k)$, then $\hat{\mathbb{G}}$ is compatible with $\mathbb{H}_k(\mathbb{G}) \simeq \mathbb{H}_k(\hat{\mathbb{G}})$.

Intuitively the precondition of part (b) should be seen as *downward compatibility*. It implies that the transition from $\mathbb{G}$ to $\hat{\mathbb{G}}$ is conservative for subgroups generated by fewer than k generators, and part (c) gives one criterion how this can be maintained while unfolding $\mathbb{G}$ at the level of larger α in non-trivial ways.

Proof. All claims follow from Observation 2.9 and Lemma 2.10 in a straightforward manner. □

The above opens up the potential for achieving successively more stringent structural conditions in an inductive fashion. Essentially the induction will be on the size $k = |\alpha|$ of the generator set of subgroups $\mathbb{G}[\alpha]$ that may form certain obstructive patterns and progresses to exclude them by replacing $\mathbb{G}$ by $\hat{\mathbb{G}}$ as in item (c) above. The crux of the matter is to reconcile the preconditions of (c), or to find suitable $\mathbb{H}_k(\mathbb{G})$ that do not spoil $\mathbb{G}[\alpha]$-compatibility. The key obstructions to be dealt with are cyclic patterns of cosets w.r.t. subgroups $\mathbb{G}[\alpha]$ to be discussed in the following section.

Some application contexts call for an analysis of symmetries of E-groups $\mathbb{G}$ that are induced by permutations of the underlying set E of generators. These are not covered by the notion of automorphisms of E-groups since those, as special homomorphisms, need to fix the generators individually (model-theoretically they are treated as constants). Similarly for E-graphs $\mathbb{H}$, automorphisms of $\mathbb{H}$ viewed as a relational structure need to respect each R_e individually, and do not account for symmetries induced by permutations of the edge colours. In both cases, permutations $\rho \in \mathrm{Sym}(E)$ induce what in model-theoretic terminology is a *renaming*, sending $\mathbb{G}$ to $\mathbb{G}^\rho$ and $\mathbb{H}$ to $\mathbb{H}^\rho$. For instance the ρ-renaming of the E-graph $\mathbb{H} = (V, (R_e)_{e \in E})$ is $\mathbb{H}^\rho = (V, (R'_e)_{e \in E})$ with $R'_{\rho(e)} = R_e$. Such a renaming reflects a *symmetry* if it leaves the underlying structure invariant up to isomorphism.

Definition 2.12. [symmetry over E]
A permutation $\rho \in \mathrm{Sym}(E)$ of the set E is a *symmetry* of an E-group $\mathbb{G}$ if the renaming of generators according to ρ yields an isomorphic E-group, $\mathbb{G}^\rho \simeq \mathbb{G}$. Similarly, ρ is a symmetry of the E-graph $\mathbb{H}$ if the renaming of its edge relations according to ρ yields an isomorphic E-graph: $\mathbb{H}^\rho \simeq \mathbb{H}$.

For instance, the trees $\mathbb{T}(E, n)$ in Biggs' construction are *fully symmetric* in the sense that every $\rho \in \mathrm{Sym}(E)$ is a symmetry; the same is then true of the resulting E-group $\mathbb{G} = \mathrm{sym}(\mathbb{T}(E, n))$ and its Cayley graph $\mathbb{CG}$.

3 Coset cycles and acyclicity criteria

The common notion of *large girth* for Cayley graphs naturally leads to There is a basic notion of n-acyclicity for E-groups that forbids non-trivial *generator cycles* (i.e. representations of $1 \in \mathbb{G}$ by reduced generator words) of lengths up to n. This account matches the graph-theoretic notion of *girth* for the associated Cayley graph, simply because the length of the shortest shortest non-trivial generator cycle in $\mathbb{G}$ is the length of the shortest graph cycle in $\mathbb{CG}$, i.e. its girth. We are interested in a more liberal notion of cycles, which leads to a more restrictive notion of acyclicity that forbids short *coset cycles*, i.e. cyclic configurations of cosets $g_i \mathbb{CG}[\alpha_i]$.

Definition 3.1. [coset cycle]
Let $\mathbb{G}$ be an E-group, $n \geqslant 2$. A *coset cycle* of length n in $\mathbb{G}$ is a cyclically indexed sequence of pointed cosets $(g_i \mathbb{G}[\alpha_i], g_i)_{i \in \mathbb{Z}_n}$ w.r.t. subgroups $\mathbb{G}[\alpha_i]$ for $\alpha_i \subseteq E$ satisfying these conditions:

(i) (connectivity) $g_{i+1} \in g_i \mathbb{G}[\alpha_i]$, i.e. $g_i \mathbb{G}[\alpha_i] = g_{i+1} \mathbb{G}[\alpha_i]$;

(ii) (separation) $g_i \mathbb{G}[\alpha_{i,i-1}] \cap g_{i+1} \mathbb{G}[\alpha_{i,i+1}] = \emptyset$,

where $\alpha_{i,j} := \alpha_i \cap \alpha_j$.

We sometimes put a focus on coset cycles whose constituent cosets stem from a restricted family of generated subgroups, and especially from $\Gamma_k(\mathbb{G})$ for some $1 \leqslant k \leqslant |E|$. With terminology like *coset cycle w.r.t.* Γ_k we then refer to coset cycles $(g_i \mathbb{G}[\alpha_i])_{i \in \mathbb{Z}_n}$ with $\alpha_i \in \Gamma_k$, i.e. with $|\alpha_i| < k$.

Definition 3.2. [N-acyclicity]
For $N \geq 2$, an E-group $\mathbb{G}$ or its Cayley graph $\mathbb{CG}$ are N*-acyclic* if they admit no coset cycles of lengths up to N.

Correspondingly N-acyclicity w.r.t. Γ_k forbids coset cycles $(g_i\mathbb{G}[\alpha_i])_{i\in\mathbb{Z}_n}$ of lengths $n \leqslant N$ with $\alpha_i \in \Gamma_k$. Obviously, non-trivial generator cycles are very special coset cycles with singleton sets $\alpha_i = \{e_i\}$. So N-acyclicity w.r.t. $\Gamma_2(\mathbb{G})$ precisely says that the girth of $\mathbb{G}$ or $\mathbb{C}\mathbb{G}$ is larger than N. Also note that N-acyclicity w.r.t. $\Gamma_k(\mathbb{G})$ in particular implies outright N-acyclicity for $\mathbb{G}[\alpha]$ for all $|\alpha| \leqslant k$. For G itself outright N-acyclicity is the same as N-acyclicity w.r.t. $\Gamma_{|E|}(\mathbb{G})$.

It is important to note that the graph-theoretic diameter of an α_i-coset in the Cayley graph or the cardinality of $\mathbb{G}[\alpha_i]$ cannot be uniformly bounded (e.g. in terms of $|\alpha_i|$). Therefore no level of generator acyclicity captures any fixed level of coset acyclicity.

From now on all reference to acyclicity will be to coset acyclicity.

The lowest level of coset acyclicity, viz. N-acyclicity for $N = 2$, is of special interest. It is easy to check that the condition for 2-acyclicity is equivalent to an intersection condition on pairs of cosets, which is reminiscent of a notion of simple connectivity.

Observation 3.3. *An* E*-group* $\mathbb{G}$ *is* 2*-acyclic if, and only if, for all* $\alpha_1, \alpha_2 \subsetneq E$,

$$\mathbb{G}[\alpha_1] \cap \mathbb{G}[\alpha_2] = \mathbb{G}[\alpha_1 \cap \alpha_2].$$

The following associates acyclicity criteria with closure properties and minimal supporting sets of generators.

Remark 3.4. *Consider an element* $g \in \mathbb{G}$ *and its* α*-component* $B = g\mathbb{G}[\alpha] \subseteq \mathbb{G}$ *for some* $\alpha \subseteq E$ *in an* E*-group* $\mathbb{G}$.

- *(i) If* $\mathbb{G}$ *is* 2*-acyclic then there is a unique* $\subseteq$*-minimal generator set* $\alpha_g \subseteq E$ *such that* $g \in \mathbb{G}[\alpha_g]$*, which is obtained as* $\alpha_g = \bigcap\{\alpha' \subseteq E \colon g \in \mathbb{G}[\alpha']\}$.
- *(ii) If* $\mathbb{G}$ *is* 3*-acyclic then there is a unique* $\subseteq$*-minimal generator set* $\alpha_B \subseteq E$ *such that* $\mathbb{G}[\alpha_B] \cap B \neq \emptyset$*, which is obtained as* $\alpha_B = \bigcap\{\alpha_h \subseteq E \colon h \in B\}$.

Proof. Observation 3.3 implies that the family of subsets $\alpha' \subseteq E$ for which $g \in \mathbb{G}[\alpha']$ is closed under intersections; this immediately implies claim (i). Towards claim (ii) consider two elements $g_i \in B$ and their supporting $\alpha_i := \alpha_{g_i}$ according to (i) for $i = 1, 2$. We need to show that also $\alpha_0 := \alpha_1 \cap \alpha_2$ supports B in the sense that $\mathbb{G}[\alpha_0] \cap B \neq \emptyset$. For this consider the potential 3-cycle of cosets $1\mathbb{G}[\alpha_1] = \mathbb{G}[\alpha_1]$, $g_1\mathbb{G}[\alpha]$ and $g_2\mathbb{G}[\alpha_2] = \mathbb{G}[\alpha_2]$. As $\mathbb{G}$ does not admit 3-cycles of cosets, at least one of the three instances of the separation conditions in Definition 3.1 must fail. We argue that each such failure yields an element in $B \cap \mathbb{G}[\alpha_0]$. Failure in the link $\mathbb{G}[\alpha_1]$ means that there is some $g' \in g_1\mathbb{G}[\alpha_1 \cap \alpha] \cap 1\mathbb{G}[\alpha_1 \cap \alpha_2]$; now $g' \in g_1\mathbb{G}[\alpha_1 \cap \alpha]$ implies that $g' \in B$ since $g_1 \in B$, and $g' \in \mathbb{G}[\alpha_0]$ as $1\mathbb{G}[\alpha_1 \cap \alpha_2] = \mathbb{G}[\alpha_0]$. Failure in the link $\mathbb{G}[\alpha_2]$ is entirely symmetric. Failure in the link $\mathbb{G}[\alpha]$ finally means that there is some $g' \in g_1\mathbb{G}[\alpha_1 \cap \alpha] \cap g_2\mathbb{G}[\alpha_2 \cap \alpha]$ which implies that $g' \in B$ as before. Moreover $g' \in \mathbb{G}[\alpha_i]$ as $g_i \in \mathbb{G}[\alpha_i]$ and $g' \in g_i\mathbb{G}[\alpha_i]$; so $g' \in \mathbb{G}[\alpha_1] \cap \mathbb{G}[\alpha_2] = \mathbb{G}[\alpha_1 \cap \alpha_2] = \mathbb{G}[\alpha_0]$. □

Lemma 3.5. *If an* E*-group* $\mathbb{G}$ *is compatible with* $\mathbb{C}\mathbb{G}[\alpha]$ *for every* $\alpha \subsetneq E$ *then* $\mathbb{G}$ *is* 3*-acyclic.*

Proof. Assume that $(g_i\mathbb{G}[\alpha_i], g_i)_{i\in\mathbb{Z}_3}$ formed a coset cycle in $\mathbb{G}$. The separation condition implies that $\alpha_i \subsetneq E$ and that, for instance,

$$(*) \qquad g_0\mathbb{G}[\alpha_{0,2}] \cap g_1\mathbb{G}[\alpha_{0,1}] = \emptyset.$$

By the connectivity condition, there are $w_j \in \alpha_j^*$ such that

$$g_0^{-1}g_1 = [w_0]_{\mathbb{G}}, \;\; g_1^{-1}g_2 = [w_1]_{\mathbb{G}}, \;\; g_2^{-1}g_0 = [w_2]_{\mathbb{G}}.$$

Clearly $[w_2w_0w_1]_{\mathbb{G}} = 1$. Let $w_{0,j} \in \alpha_{0,j}^*$ be the projection of the word w_j to α_0^*, as obtained by delition of all letters $e \notin \alpha_0$. If $\mathbb{G}$ is compatible with $\mathbb{G}[\alpha_0]$ then $[w_2w_0w_1]_{\mathbb{G}} = 1$ implies that $[w_2w_0w_1]_{\mathbb{H}} = 1 \in \mathrm{sym}(\mathbb{H})$ for $\mathbb{H} = \mathbb{C}\mathbb{G}[\alpha_0]$, which in turn implies that $[w_{0,2}w_0w_{0,1}]_{\mathbb{G}} = 1$ since $[e]_{\mathbb{H}}$ is trivial for $e \notin \alpha_0$. But the operation of the corresponding sequence of generators of $\mathrm{sym}(\mathbb{C}\mathbb{G}[\alpha_0])$ maps the element $g := [w_{0,2}]_{\mathbb{G}}^{-1}$ (via 1 and $g_0^{-1}g_1$) to $g' := g_0^{-1}g_1[w_{0,1}]_{\mathbb{G}}$. As $g \in \mathbb{G}[\alpha_{0,2}]$ and $g' \in g_0^{-1}g_1\mathbb{G}[\alpha_{0,1}]$, which by $(*)$ are disjoint subsets of $\mathbb{G}[\alpha_0]$, it follows that $[w_{0,2}w_0w_{0,1}]_{\mathbb{G}} \neq 1$, a contradiction. Analogously, $\mathbb{G}$ cannot admit 2-cycles. □

Lemma 3.6. *N-acyclicity of $\mathbb{G}[\alpha]$ is preserved under inverse homomorphisms that are injective on α'-generated subgroups for all $\alpha' \subsetneq \alpha$.*

Proof. If $h\colon \hat{\mathbb{G}}[\alpha] \to \mathbb{G}[\alpha]$ is a homomorphism of α-groups that is a local isomorphism in restriction to each $\hat{\mathbb{G}}[\alpha']$ for $\alpha' \subsetneq \alpha$, then h maps a coset cycle in $\hat{\mathbb{G}}$ to a coset cycle in $\mathbb{G}$. The connectivity condition is obviously maintained under h. The crux of the matter is the separation condition for links in a potential coset cycle. As each α_i-coset in $\hat{\mathbb{G}}$, for $\alpha_i \subsetneq \alpha$, is mapped bijectively onto its image coset in $\mathbb{G}$, so are the disjoint critical $\alpha_{i,i\pm1}$-cosets as its subsets. □

Recall Lemma 2.11, which is to be used to eliminate N-cycles of cosets with increasing numbers of generators, in an inductive treatment. More specifically, part (c) of that lemma points at the way to rule out coset cycles of length up to N formed by α-cosets for $\alpha \in \Gamma_k$ by achieving compatibility with suitable E-graph configurations $\mathbb{H}_k(\mathbb{G})$. The relevant configurations in $\mathbb{H}_k(\mathbb{G})$ will be suitable amalgams of E-graphs $\mathbb{C}\mathbb{G}[\alpha]$ for $\alpha \in \Gamma_k$, which need to be designed so as to meet the downward compatibility criterion of part (b). Towards coset acyclicity we can use amalgamation chains that unfold potential cycles.

4 Free amalgams

Free amalgams of copies of E-graphs $\mathbb{C}\mathbb{G}[\alpha_i]$ seek to superpose these copies with just those identifications that are forced by shared generator edges. We first look at superpositions of two $\mathbb{C}\mathbb{G}[\alpha_i]$, then expand the idea to certain chains or clusters of several $\mathbb{C}\mathbb{G}[\alpha_i]$. In each case, the isomorphism type of the resulting E-graphs is fully determined by the isomorphism types of the constituents $\mathbb{C}\mathbb{G}[\alpha_i]$ – and not by the manner in which the constituents are embedded in $\mathbb{C}\mathbb{G}$ as weak subgraphs $\mathbb{C}\mathbb{G}[\alpha_i] \subseteq_w \mathbb{C}\mathbb{G}$. If $\mathbb{G}$ satisfies appropriate acyclicity conditions, however, the free amalgam will be naturally isomorphic to the corresponding embedded weak subgraph of $\mathbb{CG}$ on a union of α_i-cosets of $\mathbb{G}$ (cf. Observations 4.2, 4.4, 4.8). For instance, if $\mathbb{G}$ is 2-acyclic the weak subgraphs of $\mathbb{CG}$ on overlapping cosets of the form $gG[\alpha_1] \cup gG[\alpha_2] \subseteq G$ (with induced α_i-edges on the α_i-coset) will all be isomorphic to the free amalgam $\mathbb{CG}[\alpha_1] \oplus \mathbb{CG}[\alpha_2]$ (Observation 4.2).

Definition 4.1. [free amalgam]
Let $\alpha_1, \alpha_2 \subseteq E$ with intersection $\alpha_0 = \alpha_1 \cap \alpha_2$, $\mathbb{G}$ an E-group. The *free amalgam* $\mathbb{CG}[\alpha_1] \oplus \mathbb{CG}[\alpha_2]$ of the E-graphs $\mathbb{CG}[\alpha_i] \subseteq_w \mathbb{CG}$ for $i = 1, 2$ is the E-graph $\mathfrak{A}$ obtained as the result of free amalgamation of disjoint copies of relational structures $\mathfrak{A}_i \simeq \mathbb{CG}[\alpha_i]$ via the shared embedded weak substructure $\mathfrak{A}_0 \simeq \mathbb{CG}[\alpha_0] \subseteq_w \mathbb{CG}[\alpha_i]$.

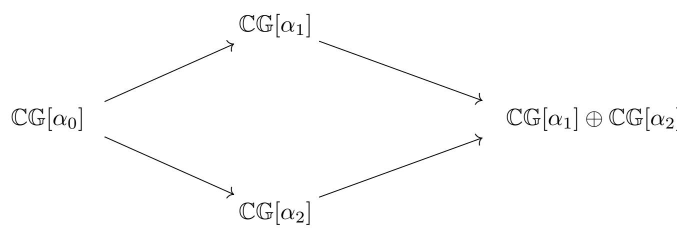

The arrows in the diagram represent injective relational homomorphisms, and the target structure interprets R_e for $e \in \alpha_1 \cup \alpha_2$ as the union of the homomorphic images of R_e from either constituent.

In related model-theoretic terminology the free amalgam of the $\mathbb{CG}[\alpha_i]$ is the *disjoint union over* $\mathbb{CG}[\alpha_0]$ of $\mathbb{CG}[\alpha_1]$ and $\mathbb{CG}[\alpha_2]$. It may be obtained from canonical disjoint isomorphic copies $\mathfrak{A}_i := \mathbb{CG}[\alpha_i] \times \{i\} \simeq \mathbb{CG}[\alpha_i]$ through identification of $a_1 = (g_1, 1) \in A_1$ with $a_2 = (g_2, 2) \in A_2$ precisely for $g_1 = g_2 \in \mathbb{G}[\alpha_0]$. That the resulting graph is a strict E-graph is seen as follows. It is an E-graph since for $e \in \alpha_0$ the overlap corresponding to $\mathbb{G}[\alpha_0]$ is closed under R_e-edges; it is strict since the amalgam does not import loops, and also cannot have multiple edges as two vertices in the overlap corresponding to $\mathbb{G}[\alpha_0]$ can be linked by at most one R_e even in $\mathbb{CG}$.

A seemingly more general concept of free amalgamation would specify an isomorphism between specific α_0-cosets in the $\mathbb{CG}[\alpha_i]$ to induce the overlap; alternatively one could specify a pair of elements from the $\mathbb{CG}[\alpha_i]$ to be identified, which in turn induces a canonical isomorphism between corresponding α_0-cosets. Such pointed copies of $\mathbb{CG}[\alpha_i]$ are a natural choice when we iterate the construction as in coset chains (see below). For the free amalgam of just two $\mathbb{CG}[\alpha_i]$ any such choices lead to isomorphic results. This is due to the internal homogeneity of Cayley graphs of E-groups.

Observation 4.2. *The map that sends an element of the free amalgam* $\mathbb{CG}[\alpha_1] \oplus \mathbb{CG}[\alpha_2]$ *to the group element in* $\mathbb{G}$ *that it stems from is a homomorphism of* E-*graphs; if the* E-*group* $\mathbb{G}$ *is* 2-*acyclic then this homomorphism is injective so that* $\mathbb{CG}[\alpha_1] \oplus \mathbb{CG}[\alpha_2]$ *is realised as a weak substructure of* $\mathbb{CG}$ *(via a canonical* *isomorphism). $\mathbb{G}$ is 2-acyclic if, and only if, $\mathbb{CG}[\alpha_1] \oplus \mathbb{CG}[\alpha_2] \subseteq_w \mathbb{CG}$ for any two $\alpha_i \subsetneq E$.*

4.1 Amalgamation chains and clusters

Iterated or multiple free amalgams of more than two $\mathbb{CG}[\alpha_i]$ can be defined in more general contexts, but special restrictions need to be in place to guarantee that he result is a strict E-graph, or even just an E-graph. We only discuss two patterns of special interest to us: *amalgamation chains* and *amalgamation clusters*. Both patterns are motivated by desirable acyclicity properties w.r.t. local overlaps between $\mathbb{G}[\alpha_i]$-cosets in $\mathbb{G}$. In other words, they unfold overlaps among a family of $\mathbb{CG}[\alpha_i]$ in the 'tree-like' pattern encountered in sufficiently acyclic $\mathbb{G}$. This is being brought out in Observations 4.4 and 4.8, respectively, which extend the basic idea of Observation 4.2 above to the two more complex patterns. While the pattern of *amalgamation chains* corresponds to a linear arrangement as in a tree branch, the pattern of *amalgamation clusters* corresponds to the branching in a single parent node.

Definition 4.3. [amalgamation chain]
Let $\mathbb{G}$ be an E-group, $N \geqslant 1$ and $\alpha_i \subsetneq E$ for $1 \leqslant i \leqslant N$ with intersections $\alpha_{i,i+1} := \alpha_i \cap \alpha_{i+1}$ for $1 \leqslant i < N$. Consider the sequence of pointed E-graphs $(\mathbb{CG}[\alpha_i], g_i)_{1 \leqslant i \leqslant N}$ and assume that the $g_i \in \mathbb{G}[\alpha_i] \subseteq_w \mathbb{CG}$ are such that the cosets $1\mathbb{G}[\alpha_{i-1,i}] \subseteq \mathbb{G}[\alpha_i]$ and $g_i\mathbb{G}[\alpha_{i,i+1}] \subseteq \mathbb{G}[\alpha_i]$ are disjoint in $\mathbb{G}[\alpha_i]$. In this situation, the *free amalgamation chain* $\bigoplus_{i=1}^{N}(\mathbb{CG}[\alpha_i], g_i)$ is the E-graph $\mathfrak{A}$ obtained as the result of simultaneous free amalgamation of disjoint copies of relational structures $\mathfrak{A}_i \simeq \mathbb{CG}[\alpha_i]$ via the shared embedded weak substructures $\mathfrak{A}_{i,i+1} \simeq g_i\mathbb{CG}[\alpha_{i,i+1}] \subseteq_w \mathbb{CG}[\alpha_i]$ and $\mathfrak{A}_{i,i+1} \simeq 1\mathbb{CG}[\alpha_{i,i+1}] \subseteq_w \mathbb{CG}[\alpha_{i+1}]$.

Note that the precondition on disjoint overlaps w.r.t. the next neighbours in the chain ensures that there is no interference between the pairwise amalgamation processes between next neighbours. Also note that (the images of) the cosets $g_i\mathbb{G}[\alpha_{i,i+1}]$ for $1 \leqslant i < N$ are separators along the chain.

Observation 4.4. *Whenever the amalgamation chain $\bigoplus(\mathbb{CG}[\alpha_i], g_i)$ is defined, there is a unique homomorphism from $\bigoplus(\mathbb{CG}[\alpha_i], g_i)$ to $\mathbb{CG}$, which maps 1 in (the isomorphic copy of) $\mathbb{CG}[\alpha_1]$ to $1 \in \mathbb{CG}$. If the E-group $\mathbb{G}$ is N-acyclic then this homomorphism is injective and $\bigoplus(\mathbb{CG}[\alpha_i], g_i)$ is realised as a weak substructure of $\mathbb{CG}$ (via an essentially canonical isomorphism).*
$\mathbb{G}$ is N-acyclic if, and only if, $\bigoplus_i(\mathbb{CG}[\alpha_i], g_i) \subseteq_w \mathbb{CG}$ for any sequence of up to N many pointed generated subgroups $(\mathbb{G}[\alpha_i], g_i)$ such that the amalgamation chain is defined.

For the contrapositive of the 'if'-part of last claim consider an amalgamation chain $\bigoplus_{i=1}^{n}(\mathbb{CG}[\alpha_i], g_i)$ that is not injectively mapped into $\mathbb{CG}$ by the natural homomorphism. For an ℓ that is minimal with the property that $1 \leqslant k < k + \ell \leqslant N$ and that the subchain $\bigoplus_{i=k}^{\ell}(\mathbb{CG}[\alpha_i], g_i)$ is not injectively embedded, its homomorphic image in $\mathbb{CG}$ constitutes a coset cycle.

Definition 4.5. [amalgamation cluster]
Let $\mathbb{G}$ be an E-group, $N \geqslant 1$. Consider a family of generated subgroups $\mathbb{G}[\alpha_i]$ for $\alpha_i \subsetneq E$ for $1 \leqslant i \leqslant N$ which are 2-acyclic with Cayley graphs $\mathbb{CG}[\alpha_i] \subseteq_w \mathbb{CG}$. For $1 \leqslant i \leqslant N$ let $\alpha_{i,i+1} := \alpha_i \cap \alpha_{i+1}$. The *free amalgamation cluster* $\bigoplus_{i=1}^N \mathbb{CG}[\alpha_i]$ is the E-graph $\mathfrak{A}$ obtained as the result of simultaneous free amalgamation of disjoint copies of relational structures $\mathfrak{A}_i \simeq \mathbb{CG}[\alpha_i]$ where elements of $\mathfrak{A}_i$ and $\mathfrak{A}_j$ are identified precisely if they correspond to the same element of the group $\mathbb{G}[\alpha_{i,j}]$.

We need to argue for existence (well-definedness) of this amalgam.

Lemma 4.6. *In the situation of Definition 4.5 the free amalgamation cluster* $\bigoplus_{i=1}^N \mathbb{CG}[\alpha_i]$ *is isomorphic to the quotient of the disjoint union of the* $\mathbb{CG}[\alpha_i]$*, as represented over the universe* $\bigcup_i (\mathbb{G}[\alpha_i] \times \{i\})$*, w.r.t. the equivalence relation* $\sim$ *that identifies* (g_i, i) *with* (g_j, j) *if, and only if,* $g_i = g_j \in \mathbb{G}[\alpha_{i,j}]$*. This quotient structure is a strict* E*-graph.*

Proof. Let $\mathfrak{A}$ be the disjoint union of the $\mathbb{CG}[\alpha_i]$ over the universe $A = \bigcup_i (\mathbb{G}[\alpha_i] \times \{i\})$. The following shows that $\sim$ is transitive, hence an equivalence relation (reflexivity and symmetry are obvious). For instance, if $(g_1, 1) \sim (g_2, 2) \sim (g_3, 3)$ where $g_i \in \mathbb{G}[\alpha_i]$, $i = 1, 2, 3$, we need to show that $(g_1, 1) \sim (g_3, 3)$. The assumption that $(g_1, 1) \sim (g_2, 2)$ implies that $g_1 = g_2 \in \mathbb{G}[\alpha_{1,2}]$ while $(g_2, 2) \sim (g_3, 3)$ implies that $g_2 = g_3 \in \mathbb{G}[\alpha_{2,3}]$. By the precondition on intersections between the $\mathbb{G}[\alpha_{i,j}]$ in the definition of the cluster (here an instance of 2-acyclicity within $\mathbb{G}[\alpha_2]$), it follows that $g_1 = g_2 = g_3 \in \mathbb{G}[\alpha_{1,2} \cap \alpha_{2,3}] \subseteq \mathbb{G}[\alpha_{1,3}]$, whence $(g_1, 1) \sim (g_3, 3)$ follows. That the quotient structure $\mathfrak{A}/\sim$ does not induce loops, edges w.r.t. multiple R_e or branching w.r.t. a single R_e follows from an analysis of when equivalence classes $[g_i, i]$ and $[g_j, j]$ are linked by an e-edge in the quotient. This is the case if, and only if, $(g_i, i) \sim (g, k)$ and $(g'_j, j) \sim (g', k)$ for some $g, g' \in \mathbb{G}[\alpha_k]$ with $e \in \alpha_k$ and $g' = ge$. It follows that $g \neq g'$ and $[g_i, i] = [g, k] \neq [g', k] = [g'_j, j]$ (no loops) and that $g' \neq ge'$ for any $e' \neq e$ (no multiple edges). If $[g, i]$ were incident with two e-edges in the quotient, say with e-edges to $[g'_j, j]$ and to $g'_k, k]$, we may w.l.o.g. assume that $e \in \alpha_j \cap \alpha_k$ and hence $e \in \alpha_{j,k}$. It follows that $(g_i, i) \sim (g_j, j) \sim (g_k, k)$ for suitable $g \in \mathbb{G}[\alpha_{j,k}]$ so that both $g'_j = ge$ and $g'_k = ge$ in $\mathbb{G}[\alpha_{j,k}]$. This shows that $g'_j = g'_k$ whence also $(g'_j, j) \sim (g'_k, k)$ so that the two e-edges are one and the same. □

Remark 4.7. *In the situation of Definition 4.5 consider an element* x *of the free amalgamation cluster* $\bigoplus_{i=1}^N \mathbb{CG}[\alpha_i]$ *and let* $M \subseteq \{1, \ldots, N\}$ *be the subset of indices* i *for which* x *is represented in the copy of* $\mathbb{CG}[\alpha_i]$ *in* $\bigoplus_{i=1}^N \mathbb{CG}[\alpha_i]$*. Then* x *determines a unique minimal set* $\alpha_x \subseteq E$ *such that* x *is represented in* $\mathbb{CG}[\alpha_x] \subseteq_w \mathbb{CG}[\alpha_i]$ *for every* $i \in M$*.*

Proof. The set α_x arises as the intersection of all those $\mathbb{CG}[\alpha'] \subseteq_w \mathbb{CG}[\alpha_i]$ which represent it in the above sense. This family of generator subsets is closed under intersection due to the free nature of the amalgamation cluster and, for reasoning inside constituents $\mathbb{CG}[\alpha_i]$, their 2-acyclicity (cf. part (i) of Remark 3.4). □

Observation 4.8. *Whenever the amalgamation cluster $\bigoplus \mathbb{C}\mathbb{G}[\alpha_i]$ is defined, there is a unique homomorphism from $\bigoplus \mathbb{C}\mathbb{G}[\alpha_i]$ to $\mathbb{C}\mathbb{G}$, which maps 1 in (the isomorphic copies) of the $\mathbb{C}\mathbb{G}[\alpha_i]$ to $1 \in \mathbb{C}\mathbb{G}$. If the* E*-group $\mathbb{G}$ is 2-acyclic then this homomorphism is injective so that $\bigoplus \mathbb{C}\mathbb{G}[\alpha_i]$ is realised as a weak substructure of $\mathbb{C}\mathbb{G}$ (via an essentially canonical isomorphism).*

4.2 Connected components of free amalgams

The structure of β-connected components of free amalgamation chains and clusters is important when these configurations occur as components of an E-graph $\mathbb{H}$ from which an E-group $\mathbb{G} = \mathrm{sym}(\mathbb{H})$ is generated as through permutation group action. Recall that the β-generated subgroup $\mathbb{G}[\beta]$ of $\mathbb{G} = \mathrm{sym}(\mathbb{H})$ is fully determined by the permutation group action on the β-connected components of $\mathbb{H}$; and hence that some β-group $\hat{\mathbb{G}}[\beta]$ admits a homomorphism to $\mathbb{G}[\beta]$ if, and only if, it is compatible with all β components of $\mathbb{H}$.

We start with the simple situation of just a binary free amalgam $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2]$ for $\alpha_1, \alpha_2 \subseteq E$ with intersection $\alpha_0 = \alpha_1 \cap \alpha_2$. In this, as in the remaining cases, we want to see that β-connected components are again free amalgams of the corresponding kind, and either trivial or built from smaller $\mathbb{C}\mathbb{G}[\alpha_i]$ (this can be regarded as a notion of downward compatibility for the corresponding class of configurations). The proofs show that acyclicity criteria for the constituent $\mathbb{C}\mathbb{G}[\alpha_i]$ are essential in this.

Lemma 4.9. *Let $\alpha_1, \alpha_2, \beta \subsetneq E$, $\mathbb{G}$ an* E*-group. If $\mathbb{G}[\alpha_1]$ and $\mathbb{G}[\alpha_2]$ are 2-acyclic, then the β-connected components of vertices in $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2]$ are either isomorphic to one of the $\mathbb{C}\mathbb{G}[\beta \cap \alpha_i]$ or to a free amalgam of the form $\mathbb{C}\mathbb{G}[\beta \cap \alpha_1] \oplus \mathbb{C}\mathbb{G}[\beta \cap \alpha_2]$.*

Proof. Let B be the vertex set of the β-component in question. If B is fully contained in (the isomorphic copy of) one of the constituent $\mathbb{C}\mathbb{G}[\alpha_i]$, then the β-component is isomorphic, as an E-graph, to $\mathbb{C}\mathbb{G}[\beta \cap \alpha_i]$ for that α_i. Otherwise B must contain a vertex v in the overlap of the constituent $\mathbb{C}\mathbb{G}[\alpha_i]$ and $\beta \cap \alpha_i \neq \emptyset$ for $i = 1, 2$. For $i = 1, 2$ let A_i be the vertex set corresponding to the constituent $\mathbb{C}\mathbb{G}[\alpha_i]$ in $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2]$. So $A_1 \cap A_2$ corresponds to $\mathbb{G}[\alpha_0] \subseteq \mathbb{G}[\alpha_i]$ where $\alpha_0 = \alpha_1 \cap \alpha_2$. Considering $B \cap A_i$ within (the isomorphic copy of) $\mathbb{C}\mathbb{G}[\alpha_i]$, the $(\beta \cap \alpha_i)$- and $(\beta \cap \alpha_0)$-cosets of the shared vertex v cannot form a 2-cycle of cosets in $\mathbb{C}\mathbb{G}[\alpha_i]$. This implies that B intersects the overlap region $A_1 \cap A_2$ in a single $(\beta \cap \alpha_0)$-coset. This further implies that B extends from this copy of $\mathbb{C}\mathbb{G}[\beta \cap \alpha_1 \cap \alpha_2]$ to a full copy of $\mathbb{C}\mathbb{G}[\beta \cap \alpha_i]$ within A_i for $i = 1, 2$. So the β-component in question is isomorphic to the free amalgam $\mathbb{C}\mathbb{G}[\beta \cap \alpha_1] \oplus \mathbb{C}\mathbb{G}[\beta \cap \alpha_2]$ as claimed. $\square$

Lemma 4.10. *Let $\alpha_i \subsetneq E$ for $1 \leqslant i \leqslant N$, $\beta \subseteq E$, and let the* E*-group $\mathbb{G}$ and elements $g_i \in \mathbb{G}[\alpha_i]$ be such that the free amalgamation chain $\bigoplus_{i=1}^{N} (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ is defined. If the $\mathbb{G}[\alpha_i]$ are 2-acyclic, then the β-connected components of vertices in the* E*-graph $\bigoplus_{i=1}^{N} (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ are isomorphic to free amalgamation chains of* *the form* $\bigoplus_{i=s}^{t}(\mathbb{CG}[\beta_i], h_i)$ *for* $\beta_i := \beta \cap \alpha_i$*, some* $1 \leqslant s \leqslant t \leqslant N$*, and suitable choices of* $h_i \in \mathbb{G}[\alpha_i]$*.*

Proof. This follows along the same lines as the binary case in Lemma 4.9. Let B be the β-component under consideration and choose the index s minimal such that the $\mathbb{CG}[\alpha_s]$-copy in the chain contributes to B. Starting from a vertex in B from the $\mathbb{CG}[\alpha_s]$-copy we may proceed through one overlap at a time until the sub-chain may die out. The h_i are chosen in the (copies of the) overlaps $\mathbb{CG}[\alpha_{i,i+1}]$ (where $\alpha_{i,i+1} = \alpha_i \cap \alpha_{i+1}$) that contribute to B. The condition of 2-acyclicity for the $\mathbb{G}[\alpha_i]$ guarantees that the contribution of each overlap $\mathbb{CG}[\alpha_{i,i+1}]$ is a single $(\beta \cap \alpha_{i,i+1})$-coset (note that $\beta \cap \alpha_{i,i+1} = \beta_i \cap \beta_j$). ☐

Definition 4.11. [cluster property]
An E-group $\mathbb{G}$ has the *cluster property* if for all free amalgamation clusters $\bigoplus_{i \in I} \mathbb{CG}[\alpha_i]$ with $\alpha_i \subsetneq E$ every β-connected component B for $\beta \subsetneq E$ contains an element x such that $\alpha_x = \alpha_B := \bigcap\{\alpha_y : y \in B\}$ and, as a β-graph, is isomorphic to the free amalgamation cluster $\bigoplus_{i \in M} \mathbb{CG}[\beta_i]$ where $\beta_i = \beta \cap \alpha_i$ and $M = \{i \in I : \beta_i \neq \emptyset\}$.

Lemma 4.12. *For any* E*-group* $\mathbb{G}$ *and* $\alpha \subseteq E$*, if* $\mathbb{G}[\alpha']$ *is* 2*-acyclic and has the cluster property for all* $\alpha' \subsetneq \alpha$ *then* $\mathbb{G}[\alpha]$ *has the cluster property.*

Proof. Let $\beta \subsetneq \alpha$ and consider a β-component B of $\bigoplus \mathbb{CG}[\alpha_i]$, $\beta_i := \beta \cap \alpha_i$, $M := \{i : \beta_i \neq \emptyset\}$ as in the statement. The claim is trivial if M is empty or a singleton set, as then B is a singleton or contained in a single $\mathbb{CG}[\alpha_i]$ which has the cluster property by assumption.

If M has a least two elements, say $i = 1, 2$, consider an element $x \in B$ that is an element of the constituent $\mathbb{CG}[\alpha_i]$-copies for $i = 1, 2$, neither of which covers all of B. Consider the minimal supporting set of generators $\alpha_x \subseteq \alpha_0 := \alpha_1 \cap \alpha_2$ according to Remark 4.7, and the β_i-connected components B_i of x in the constituent $\mathbb{CG}[\alpha_i]$-copies, which are β_i-cosets. The cluster property for the $\mathbb{G}[\alpha_i]$ provides $x_i \in B_i$ with $\alpha_{x_i} = \alpha_{B_i} = \bigcap\{\alpha_y : y \in B_i\}$. By minimality of α_{B_i}, $\alpha_{x_i} = \alpha_{B_i} \subseteq \alpha_x$. Now x_i is linked to x by a β_i-walk, and also (via 1 in the copy of $\mathbb{CG}[\alpha_i]$) by the composition of an α_{x_i}-walk and an α_x-walk, which is an α_0-walk (as $\alpha_{x_i} = \alpha_{B_i} \subseteq \alpha_x \subseteq \alpha_0$). Since $\mathbb{G}[\alpha_i]$ is 2-acyclic, x_i must be linked to x by a β_0-walk for $\beta_0 := \beta \cap \alpha_0$. It follows that both x_i are elements of the shared copy of $\mathbb{CG}[\alpha_0]$ and of the β-component B_0 of x within this copy. As $\mathbb{CG}[\alpha_0]$ has the cluster property, this β-component B_0 of x in $\mathbb{CG}[\alpha_0]$ has some element $x_0 \in B_0$ such that $\alpha_{x_0} = \alpha_{B_0} = \bigcap_{y \in B_0} \alpha_y$. By minimality of α_{B_0} in the shared copy of $\mathbb{CG}[\alpha_0]$, which represents x and all the x_i, we find that $\alpha_{B_0} \subseteq \alpha_{x_1} \cap \alpha_{x_2} = \alpha_{B_1} \cap \alpha_{B_2}$. Minimality of the α_{B_i} in their $\mathbb{CG}[\alpha_i]$-copies implies that $\alpha_{B_0} = \alpha_{B_1} = \alpha_{B_2}$. As α_1 and α_2 were chosen arbitrarily from among the α_i with $\beta_i \neq \emptyset$ we find that all contributing β_i-cosets in their respective $\mathbb{CG}[\alpha_i]$-copies overlap in a shared β_0-coset that is part of a shared constituent $\mathbb{CG}[\alpha_0] \subseteq_w \mathbb{CG}[\alpha_i]$ where $\alpha_0 = \bigcap_{i \in M} \alpha_i$ and $\beta_0 = \beta \cap \alpha_0$. Hence the β-reduct of the substructure $\bigoplus \mathbb{CG}[\alpha_i] \restriction B$ is isomorphic to the free amalgam $\bigoplus_{i \in M} \mathbb{CG}[\beta_i]$ whose core is the $\mathbb{CG}[\alpha_0]$-copy shared by all the B_i for $i \in M$. ☐

Corollary 4.13. *Any* E*-group* $\mathbb{G}$ *that is* 2*-acyclic has the cluster property.*

Proof. The cluster property for $\mathbb{G}[\alpha]$ is trivial for $|\alpha| \leqslant 2$ as connected components for $\beta \subsetneq \alpha$ can only be singletons or individual β-edges linking two vertices x_i whose minimal supporting sets α_{x_i} can at most differ by the label of that β-edge. An inductive application of the lemma yields the cluster property for all $\mathbb{G}[\alpha]$ and hence for $\mathbb{G} = \mathbb{G}[E]$. □

5 Construction of N-acyclic E-groups

The actual construction of finite N-acyclic E-groups $\mathbb{G}$ is achieved in an inductive process that produces, for fixed $N \geqslant 2$, a sequence of E-groups $\mathbb{G}_k$, for increasing values of k in the range between 1 and $|E|$ such that

$$\mathbb{G}_k \text{ is } N\text{-acyclic w.r.t. } \Gamma_k(\mathbb{G}_k),$$

i.e. $\mathbb{G}_k$ admits no coset cycles of length up to N with constituent cosets w.r.t. subgroups from $\Gamma_k(\mathbb{G}_k)$ (α-cosets for $|\alpha| < k$). In this inductive sequence $\mathbb{G}_{k+1}$ is obtained as $\mathbb{G}_{k+1} := \mathrm{sym}(\mathbb{H}_k)$ from an E-graph $\mathbb{H}_k$ which in turn is defined from $\mathbb{C}\mathbb{G}_k$; the E-graph $\mathbb{H}_k$ is chosen such that

– $\mathbb{G}_{k+1} \succcurlyeq \mathbb{G}_k$,

– $\Gamma_{k+1}(\mathbb{G}_{k+1}) = \Gamma_{k+1}(\mathbb{G}_k)$, i.e. $\mathbb{G}_{k+1}[\alpha] \simeq \mathbb{G}_k[\alpha]$ for $|\alpha| \leqslant k$.

The following can then be used for the passage from $\mathbb{G}_k$ to $\mathbb{G}_{k+1}$ at the level of generated subgroups $\Gamma_{k+1}(\mathbb{G}_{k+1})$ in relation to chains over $\Gamma_{k+1}(\mathbb{G}_{k+1}) = \Gamma_{k+1}(\mathbb{G}_k)$.

Lemma 5.1. *For* $n \geqslant 1$*, let* $\mathbb{G}[\alpha]$ *be compatible with all free amalgamation chains of length up to* n *with constituents* $\mathbb{C}\mathbb{G}[\alpha']$ *for* $\alpha' \subsetneq \alpha$*. Then* $\mathbb{G}[\alpha]$ *is* N*-acyclic for* $N = n + 2$*.*

Proof. The gist of the matter is that $\mathbb{G} \succcurlyeq \mathrm{sym}(\mathbb{H})$ for every free amalgamation chain $\mathbb{H}$ of length up to n. These free amalgamation chains unfold potential coset cycles. This rules out corresponding cycles in $\mathrm{sym}(\mathbb{H})$ and hence in $\mathbb{G}$ as follows. Suppose the pointed cosets $(g_i\mathbb{G}[\alpha_i], g_i)_{i\in\mathbb{Z}_N}$ formed a coset cycle in $\mathbb{G}$. Similar to the argument in Lemma 3.5, we think of cutting the cycle (this time in g_0) and test the permutation group action on a chain formed by the remaining links, viz. on the free amalgamation chain

$$\mathbb{H} = \bigoplus_{i=1}^{n} (\mathbb{C}\mathbb{G}[\alpha_i], g_i^{-1}g_{i+1})$$

of length $n = N - 2$. As before let $\alpha_{i,j} := \alpha_i \cap \alpha_j$, and let $w_i \in \alpha_i^*$ be such that $[w_i]_{\mathbb{G}} = g_i^{-1}g_{i+1}$. For the links from and to g_0 in the cycle, $w_0 \in \alpha_0^*$ and $w_{n+1} \in \alpha_{n+1}^*$, we also look at the projections of to the neighbouring constituents in the chain $\mathbb{H}$. Let $w_{0,1} \in \alpha_{0,1}^*$ be the projection of w_0 to α_1 and $w_{n,n+1} \in \alpha_{n,n+1}^*$ the projection of w_{n+1} to α_n. Let v be the element of the chain $\mathbb{H}$ that corresponds to $[w_{0,1}]_{\mathbb{G}}^{-1}$ in its first constituent $\mathbb{C}\mathbb{G}[\alpha_1]$, and consider the permutation group action of $[\prod_{i=0}^{n+1} w_i]_{\mathbb{H}} \in \mathrm{sym}(\mathbb{H})$, corresponding to $\prod_{i=0}^{n+1}(g_i^{-1}g_{i+1}) = 1 \in \mathbb{G}$, on v in $\mathbb{H}$. By the separation condition for the α_0-link of the cycle, the generator sequence w_0 has the same effect on v as its projection $w_{0,1}$ and maps v to the element corresponding to 1 in the first constituent $\mathbb{C}\mathbb{G}[\alpha_1]$ of $\mathbb{H}$; the separation condition for the α_i-links up to $i = n$ imply that $[\prod_{i=0}^{n} w_i]_{\mathbb{H}}$ maps v to the element corresponding to $g_n^{-1}g_{n+1}$ in the last constituent $\mathbb{C}\mathbb{G}[\alpha_n]$ of $\mathbb{H}$; and the separation condition for the α_{n+1}-link of the cycle shows that the final image of v is an element in the $\alpha_{n,n+1}$-component of $g_n^{-1}g_{n+1}$ in that $\mathbb{C}\mathbb{G}[\alpha_n]$ constituent of $\mathbb{H}$. But this image is necessarily distinct from v, contradicting compatibility of $\mathbb{G}$ with $\mathbb{H}$, as $[\prod_{i=0}^{n+1} w_i]_{\mathbb{G}} = 1$. □

The following lemma focuses on the underlying induction scheme in the proposed construction.

Lemma 5.2. *For any finite* E*-group* $\mathbb{G}$ *and* $n \geqslant 1$ *there is a sequence of finite* E*-groups* $(\mathbb{G}_k)_{k \leqslant |E|}$ *starting with* $\mathbb{G}_0 := \mathbb{G}$ *such that for* $k < |E|$*:*

(i) $\mathbb{G}_k \preccurlyeq \mathbb{G}_{k+1}$*,*

(ii) $\Gamma_{k+1}(\mathbb{G}_k) = \Gamma_{k+1}(\mathbb{G}_{k+1})$*,*

(iii) $\mathbb{G}_{k+1}$ *is compatible with all free amalgamation chains of length up to* n *over* $\Gamma_{k+1}(\mathbb{G}_{k+1}) = \Gamma_{k+1}(\mathbb{G}_k)$*.*

It follows that, for $N = n+2$*,* $\mathbb{G}_{k+1}$ *is* N*-acyclic w.r.t.* $\Gamma_{k+1}(\mathbb{G}_{k+1}) = \Gamma_{k+1}(\mathbb{G}_k)$*. In particular,* $\mathbb{G}_k[\alpha]$ *is* N*-acyclic for all* α *of size* $|\alpha| \leqslant k$*, and for* $\hat{\mathbb{G}} := \mathbb{G}_{|E|}$*,*

$$\mathbb{G} \preccurlyeq \hat{\mathbb{G}} \quad \textit{where } \hat{\mathbb{G}} \textit{ is } N\textit{-acyclic.}$$

Proof. The compatibility/acyclicity requirement in (iii) is vacuous at the level of $k = 0$, as is the conservative behaviour in (ii) for 1-generator subgroups, which are isomorphic to $\mathbb{Z}_2$ in any E-group. For $k \geqslant 0$, $\mathbb{G}_{k+1}$ can be obtained as $\mathbb{G}_{k+1} = \mathrm{sym}(\mathbb{H}_k)$ where $\mathbb{H}_k$ is the disjoint union of copies of $\mathbb{C}\mathbb{G}_k$ and every free amalgamation chain of length up to n with constituents $\mathbb{G}_k[\alpha_i]$ for $\alpha_i \in \Gamma_{k+1}$. Lemma 2.11 and Lemma 4.10 show that $\mathbb{G}_{k+1}$ is as required in (i) and (ii) in relation to $\mathbb{G}_k$. Condition (iii) follows by choice of $\mathbb{H}_k$ (for compatibility) and Lemma 5.1 (for acyclicity). □

Recall from Definition 2.12 the notion of symmetries for E-groups and E-graphs that are induced by permutations of the set E. It is clear from the construction steps in Lemma 5.2 above that they do not break any such symmetry in the passage from $\mathbb{G}_k$ to $\mathbb{G}_{k+1}$. It follows that the passage from $\mathbb{G}$ to an N-acyclic $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$ preserves all symmetries of $\mathbb{G}$. We state this additional feature in the theorem which otherwise just sums up the outcome of the lemma.

Theorem 5.3. *For every finite* E*-group* $\mathbb{G}$ *and* $N \geqslant 2$ *there is a finite* E*-group* $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$ *that is* N*-acyclic and fully symmetric over* $\mathbb{G}$ *in the sense that every permutation of the generator set* E *that is a symmetry of* $\mathbb{G}$ *extends to a symmetry of* $\hat{\mathbb{G}}$*:* $\mathbb{G}^\rho \simeq \mathbb{G} \Rightarrow \hat{\mathbb{G}}^\rho \simeq \hat{\mathbb{G}}$*.*

In particular, we obtain finite N-acyclic E-groups $\hat{\mathbb{G}}$ that are fully symmetric, admitting every permutation of E as a symmetry, if we start from the fully symmetric $\mathbb{G} := \mathrm{sym}(\mathbb{H})$ for the hypercube $\mathbb{H} = 2^E$.

6 Constraints on generator sequences

The building blocks of plain coset cycles are generated subgroups of the form $\mathbb{G}[\alpha] \subseteq \mathbb{G}$, which may also be seen as the images of $\alpha^* \subseteq E^*$ under the natural homomorphism

$$\begin{array}{rcl} [\;]_{\mathbb{G}} \colon E^* & \longrightarrow & \mathbb{G} \\ w = e_1 \cdots e_n & \longmapsto & [w]_{\mathbb{G}} := \prod_{i=1}^n e_i = e_1 \cdots e_n \end{array}$$

that associates a group element with any (reduced) word over E. This association naturally translates to cosets $g\mathbb{G}[\alpha]$. Alternatively, $\mathbb{G}[\alpha]$ and $g\mathbb{G}[\alpha]$ may be regarded as the α-connected components of 1 or g in the Cayley graph $\mathbb{CG}$ of $\mathbb{G}$.

A natural way of putting extra constraints on these weak subgraphs, with reasonable closure properties in terms of generator sets $\alpha \subseteq E$, is the following. Consider a fixed E-graph $\mathbb{I} = (S, (R_e)_{e \in E})$ on vertex set S. We regard $\mathbb{I}$ as a template for systematic restrictions on patterns of generator sequences, calling it a *constraint graph*.

Proviso 6.1. *We fix an* E*-graph* $\mathbb{I} = (S, (R_e)_{e \in E})$ *as a* constraint graph*, and consider only* E*-groups* $\mathbb{G}$ *that are compatible with* $\mathbb{I}$*, i.e. with* $\mathbb{G} \succcurlyeq \mathrm{sym}(\mathbb{I})$*.*

Remark 6.2. *The restriction to* E*-groups that are compatible with* $\mathbb{I}$*,* $\mathbb{G} \succcurlyeq \mathrm{sym}(\mathbb{I})$*, does imply that there is a well-defined group action of* $\mathbb{G}$ *on* $\mathbb{I}$*. But w.r.t. this group action, any* $s \in S$ *that is not incident with an* e*-edge is a fixed point of* $\pi_e \in \mathrm{sym}(\mathbb{I})$*. As* $\mathbb{I}$ *will typically not be a complete* E*-graph, the Cayley graph of* $\mathbb{CG}$ *does not map homomorphically to* $\mathbb{I}$*. But* α*-walks from* $s \in \mathbb{I}$ *do have unique lifts to* α*-walks from any* $g \in \mathbb{CG}$*. Compatibility says that lifts at the same* $g \in \mathbb{CG}$ *of different walks from* $s \in S$ *can only meet in* $\mathbb{CG}$ *above positions in which the given walks meet in* $\mathbb{I}$*.*

The idea is to regard $\mathbb{I}$ as a template for edge patterns of walks and correspondingly restricted notions of reachability and connected components. Recall from Definition 2.7 the notions of α-walks in E-graphs, which we shall now refine in connection with the constraint graph $\mathbb{I}$. According to Definition 2.7, a *walk* in an E-graph $\mathbb{H}$ is a sequence of vertices and edge labels of the form

$$s_0, e_1, s_1, e_2, \ldots, s_{n-1}, e_n, s_n$$

with $s_i \in S$ and $e_i \in E$ where $(s_i, s_{i+1}) \in R_{e_{i+1}}$ for $0 \leqslant i < n$ for some $n \in \mathbb{N}$. The edges $(s_i, s_{i+1}) \in R_{e_{i+1}}$ are the edges traversed by this walk. The above walk is a walk of length n from the source $s = s_0$ to the target $t = s_n$ and its *edge label sequence* is the word $w = e_1 \cdots e_n \in E^*$. We also say that this word w *labels a walk* (of length n, from s_0 to s_n) in $\mathbb{H}$, and describe the walk in question as a w-walk, or as an X-walk for some language $X \subseteq E^*$ if $w \in X$. In the case of $X = \alpha^*$ we also speak of α-walks instead of α^*-walks.

Note that a word $w \in E^*$ can label at most one walk from a given vertex v in any E-graph $\mathbb{H}$. That all $w \in E^*$ label walks from all $v \in \mathbb{H}$ is equivalent to $\mathbb{H}$ being complete (i.e. to each R_e being a full matching).

Definition 6.3. [$\mathbb{I}$-words and $\mathbb{I}$-walks]
For $\alpha \subseteq E$ and $s \in S$ let $\mathbb{I}[\alpha, s] \subseteq_w \mathbb{I}$ be the weak substructure of $\mathbb{I} = (S, (R_e)_{e \in E})$ whose universe is the α-connected component of s (i.e. the connected component of s in the α-reduct $\mathbb{I}{\restriction}\alpha$), and with induced R_e for all $e \in \alpha$. Natural sets of E- or α-words that occur as edge label sequences along walks on the constraint graph $\mathbb{I}$ are defined as follows:

– $\alpha^*[\mathbb{I}] \subseteq E^*$ consists of those $w \in \alpha^*$ that label a walk in $\mathbb{I}$;
– $\alpha^*[\mathbb{I}, s] \subseteq E^*$ consists of those $w \in \alpha^*$ that label a walk from s in $\mathbb{I}$;
– $\alpha^*[\mathbb{I}, s, t] \subseteq E^*$ consists of those $w \in \alpha^*$ that label a walk from s to t in $\mathbb{I}$.

Definition 6.4. [$\mathbb{I}[\alpha, s]$-component]
For an E-graph $\mathbb{H}$, $\alpha \subseteq E$ and $s \in S$, the $\mathbb{I}[\alpha, s]$-*component* of a vertex $v \in \mathbb{H}$ is the following weak substructure $\mathbb{H}[\mathbb{I}, \alpha, s; v] \subseteq_w \mathbb{H}$: its vertex set consists of those vertices that are reachable on an $\alpha^*[\mathbb{I}, s]$-walk from v (i.e. on a walk in $\mathbb{H}$ whose edge label sequence is in $\alpha^*[\mathbb{I}, s]$); its edge relations comprise those R_e-edges for $e \in \alpha$ that are traversed by $\alpha^*[\mathbb{I}, s]$-walks from v in $\mathbb{H}$.

When we speak of an $\mathbb{I}[\alpha]$-component we mean an $\mathbb{I}[\alpha, s]$-component for some $s \in S$, which is left unspecified.

For later use we define a direct product of $\mathbb{I}$ with the Cayley graph of an E-group, which reflects $\mathbb{I}$-reachability, as follows.

Definition 6.5. [direct product]
Let $\mathbb{G}$ be an E-group that is compatible with the constraint graph $\mathbb{I} = (S, E)$, $\mathbb{CG}$ the Cayley graph of $\mathbb{G}$. Then the *direct product* $\mathbb{I} \otimes \mathbb{CG}$ is the E-graph

$$\mathbb{I} \otimes \mathbb{CG} = (V, (R_e)_{e \in E})$$

with vertex set $V = S \times G$ and edge relations

$$R_e = \{((s, g), (s', g')) \colon (s, s') \in R_e \text{ in } \mathbb{I} \text{ and } g' = ge \text{ in } \mathbb{G}\} \quad \text{for } e \in E.$$

Note that all α-walks in $\mathbb{I} \otimes \mathbb{CG}$, by definition of the edge relations, trace the lifts of α-walks in $\mathbb{I}$. For the following also compare Remark 6.2 above on lifts of walks from $\mathbb{I}$ to $\mathbb{CG}$.

Observation 6.6. *Computability of $\mathbb{G}$ with $\mathbb{I}$ implies that for any (s, g) and (s', g) in the same connected component of $\mathbb{I} \otimes \mathbb{CG}$ we must have $s = s'$. It follows that connected components of $\mathbb{I} \otimes \mathbb{CG}$ are isomorphic to weak subgraphs of $\mathbb{CG}$, and that the connected components of $\mathbb{I} \otimes \mathbb{CG}$ reflect $\mathbb{I}$-reachability in $\mathbb{CG}$ in the sense that (s, g) and (t, g') are in the same α-connected component of $\mathbb{I} \otimes \mathbb{CG}$ if, and only if, $g' \in \mathbb{CG}[\mathbb{I}, \alpha, s; g]$.*

We turn to cosets, coset cycles and acyclicity criteria relative to the given constraint graph $\mathbb{I}$.

Definition 6.7. [$\mathbb{I}$-coset]
In an E-group $\mathbb{G}$ that is compatible with $\mathbb{I}$ the $\mathbb{I}[\alpha, s]$-*coset* of $g \in \mathbb{G}$ is the $\mathbb{I}[\alpha, s]$-component $\mathbb{CG}[\mathbb{I}, \alpha, s; g] \subseteq_w \mathbb{CG}$ of g.

We drop mention of g for cosets at $g = 1$, writing e.g. just $\mathbb{CG}[\mathbb{I}, \alpha, s] \subseteq_w \mathbb{CG}$ for $\mathbb{CG}[\mathbb{I}, \alpha, s; 1] \subseteq_w \mathbb{CG}$. For the following compare Definition 3.1 for plain coset cycles.

Definition 6.8. [$\mathbb{I}$-coset cycle]
Let $\mathbb{G}$ be an E-group, $n \geqslant 2$. An $\mathbb{I}$-*coset cycle* of length $n \geqslant 2$ in $\mathbb{G}$ or $\mathbb{CG}$ is a cyclically indexed sequence $(\mathbb{CG}[\mathbb{I}, \alpha_i, s_i; g_i], g_i)_{i \in \mathbb{Z}_n}$ of pointed $\mathbb{I}[\alpha_i]$-cosets for $\alpha_i \subsetneq E$, $s_i \in S$ and $g_i \in G$ satisfying these conditions:

(i) (connectivity) there is an $\alpha_i^*[\mathbb{I}, s_i, s_{i+1}]$-walk from g_i to g_{i+1}, [5]

(ii) (separation) $\mathbb{CG}[\mathbb{I}, \alpha_{i,i-1}, s_i; g_i] \cap \mathbb{CG}[\mathbb{I}, \alpha_{i,i+1}, s_{i+1}; g_{i+1}] = \emptyset$,

where $\alpha_{i,j} := \alpha_i \cap \alpha_j$.

Definition 6.9. [N-acyclicity over $\mathbb{I}$]
An E-group that is compatible with $\mathbb{I}$ is called N-*acyclic over* $\mathbb{I}$ if it does not admit any $\mathbb{I}$-coset cycles of length up to N.

The above definitions generalise corresponding definitions in the unconstrained setting. Those definitions, Definitions 3.1 and 3.2, are comprised in the above as special cases for the trivial constraint graph having a single vertex with loops for all $e \in E$.

As in the unconstrained case, 2-acyclicity over $\mathbb{I}$ is akin to a notion of simple connectivity, being equivalent to the requirement that, for all α_i and s,

$$\mathbb{CG}[\mathbb{I}, \alpha_1, s] \cap \mathbb{CG}[\mathbb{I}, \alpha_2, s] = \mathbb{CG}[\mathbb{I}, \alpha_1 \cap \alpha_2, s].$$

7 Construction of N-acyclic E-groups over I

In order to boost degrees of plain coset acyclicity to acyclicity w.r.t. $\mathbb{I}$-cosets we aim to employ compatibility with unfoldings of potential $\mathbb{I}$-coset cycles in a manner similar to the treatment of Section 5. The following lays out a strong criterion that allows for the unfolding of potential $\mathbb{I}$-coset cycles with constituents $\mathbb{CG}[\mathbb{I}, \alpha_i, s_i, g_i] \subseteq_w \mathbb{CG}[\alpha_i]$ through the unfolding of the surrounding plain cosets $\mathbb{CG}[\alpha_i]$ into free amalgamation chains. This is the same unfolding into amalgamation chains as discussed in Section 4 and used towards plain coset acyclicity in Section 5. The crucial challenge is to safeguard this unfolding and the corresponding compatibility arguments against damage through unwanted shortcuts by α-walks that do not correspond to $\alpha^*[\mathbb{I}]$-walks. It is to this end that we consider the relationship between $\alpha^*[\mathbb{I}]$-components and the halo around them that is formed by overlapping α'-connected components for $\alpha' \subsetneq \alpha$.

[5] Equivalently, $g_{i+1} \in \mathbb{CG}[\mathbb{I}, \alpha_i, s_i; g_i]$, or $\mathbb{CG}[\mathbb{I}, \alpha_i, s_i; g_i] = \mathbb{CG}[\mathbb{I}, \alpha_i, s_{i+1}; g_{i+1}]$

Definition 7.1. [$\mathbb{I}$-skeleton]
For $\alpha \subseteq E$ an $\mathbb{I}[\alpha, s]$-*skeleton* is an E-graph $\mathbb{H}$ that admits a surjective homomorphism $h\colon \mathbb{H} \to \mathbb{I}[\alpha, s]$ onto an α-connected component $\mathbb{I}[\alpha, s]$ of $\mathbb{I}$ with the following lifting property: whenever $h(v) \in S$ is incident with an e-edge in $\mathbb{I}[\alpha, s]$ then so is v in $\mathbb{H}$.
An $\mathbb{I}[\alpha]$-skeleton is a disjoint union of $\mathbb{I}[\alpha, s]$-skeletons whose homomorphic image in $\mathbb{I}$ covers all of $\mathbb{I}\restriction\alpha$ (i.e. there is at least one component for each α-connected component $\mathbb{I}[\alpha, s] \subseteq_w \mathbb{I}$.

Note that the lifting property guarantees that every $\alpha^*[\mathbb{I}]$-walk from s' in $\mathbb{I}[\alpha, s]$ has a unique lift to an $\alpha^*[\mathbb{I}]$-walk from v for every v in the pre-image of s'. In alternative terminology, the homomorphism establishes $\mathbb{H}$ as an unbranched bisimilar cover of $\mathbb{I}[\alpha, s]$ (or of $\mathbb{I}\restriction\alpha$).

Example 7.2. Every $\mathbb{CG}[\mathbb{I}, \alpha, s; g] \subseteq_w \mathbb{CG}[\alpha] \subseteq_w \mathbb{CG}$ is an $\mathbb{I}[\alpha, s]$-skeleton, via the unique homomorphism h that maps $g \in \mathbb{G}$ to $s \in S$. This map is well-defined since $\mathbb{G}$ is assumed compatible with $\mathbb{I}$ (cf. Remark 6.2 and comments after Definition 6.5).

The following definition captures the notion that these embedded $\mathbb{I}[\alpha, s]$-skeletons in $\mathbb{CG}$ are not tied by incidental overlaps between α'-cosets that are not matched by overlaps between the corresponding $\mathbb{I}[\alpha']$-cosets within the skeleton, for smaller generator sets $\alpha' \subsetneq \alpha$. Similar to the notation Γ_k for the set of all generator subsets of size less than k, we now write

$$\Gamma(\alpha) := \{\alpha' \subseteq E \colon \alpha' \subsetneq \alpha\}$$

for the collection of the critical generator subsets under consideration. Again we abbreviate corresponding families of generated subgroups and their Cayley graphs from a given E-group $\mathbb{G}$ as

$$\Gamma(\mathbb{G}[\alpha]) := (\mathbb{G}[\alpha'] \colon \alpha' \in \Gamma(\alpha))$$
$$\Gamma(\mathbb{CG}[\alpha]) := (\mathbb{CG}[\alpha'] \colon \alpha' \in \Gamma(\alpha))$$

Definition 7.3. [freeness]
The embedded $\mathbb{I}[\alpha, s]$-skeleton $\mathbb{CG}[\mathbb{I}, \alpha, s; g] \subseteq_w \mathbb{CG}$ is *free* in $g\mathbb{CG}[\alpha]$ (or in $\mathbb{CG}$) if any two cosets $g_1\mathbb{G}[\alpha_i]$ and $g_2\mathbb{G}[\alpha_2]$, for $g_i \in \mathbb{CG}[\mathbb{I}, \alpha, s; g]$ and $\alpha_i \in \Gamma(\alpha)$ that overlap in $\mathbb{CG}$ also overlap within the skeleton $\mathbb{CG}[\mathbb{I}, \alpha, s; g]$:

$$g_1\mathbb{G}[\alpha_1] \cap g_2\mathbb{G}[\alpha_2] \neq \emptyset \quad \Rightarrow \quad \mathbb{CG}[\mathbb{I}, \alpha_1, s_1; g_1] \cap \mathbb{CG}[\mathbb{I}, \alpha_2, s_2; g_2] \neq \emptyset$$

where $s_i = h(g_i)$ for the canonical homomorphism $h\colon \mathbb{CG}[\mathbb{I}, \alpha, s; g] \to \mathbb{I}$ from Example 7.2. The E-group $\mathbb{G}$ is *free over* $\mathbb{I}$ if every embedded $\mathbb{I}[\alpha, s]$-skeleton $\mathbb{CG}[\mathbb{I}, \alpha, s; g] \subseteq_w \mathbb{CG}$ for $\alpha \subseteq E$ is free in $\mathbb{CG}$.

For generated subgroups we similarly say that $\mathbb{G}[\alpha]$ *is free over* $\mathbb{I}$ if every embedded $\mathbb{I}[\alpha', s]$-skeleton $\mathbb{CG}[\mathbb{I}, \alpha', s; g] \subseteq_w \mathbb{CG}[\alpha]$ for $\alpha' \subseteq \alpha$ is free in $\mathbb{CG}[\alpha]$.
Already the special case of $\alpha_1 = \alpha_2 \in \Gamma(\alpha)$ gives an indication of the strength of this freeness condition. For $\alpha' = \alpha_1 = \alpha_2 \in \Gamma(\alpha)$ the freeness requirement

for $\mathbb{CG}[\mathbb{I},\alpha,s;g] \subseteq_w \mathbb{CG}$ in $g\mathbb{CG}[\alpha]$ says that α'-connected components of the skeleton $\mathbb{CG}[\mathbb{I},\alpha,s;g]$ are induced by α'-components of the surrounding $g\mathbb{CG}[\alpha]$. In other words, α'-reachability inside the embedded $\mathbb{I}[\alpha,s]$-skeleton agrees with α'-reachability in the surrounding $\mathbb{CG}$.

The following is an analogue of Lemma 3.6 for freeness.

Lemma 7.4. *If $\mathbb{G}[\alpha]$ is 2-acyclic, then freeness of $\mathbb{G}[\alpha]$ over $\mathbb{I}$ is preserved under inverse homomorphisms that are injective on α'-generated subgroups for $\alpha' \subsetneq \alpha$.*

Proof. Let $h\colon \hat{\mathbb{G}} \to \mathbb{G}$ be a homomorphism of E-groups (both assumed to be compatible with $\mathbb{I}$: Proviso 6.1), its restrictions to members of the family $\Gamma(\hat{\mathbb{G}}[\alpha])$ injective. The image of an embedded $\mathbb{I}[\alpha,s]$-skeleton in $\mathbb{C}\hat{\mathbb{G}}$ under h is an embedded $\mathbb{I}[\alpha,s]$-skeleton in $\mathbb{CG}$. Assume that an $\mathbb{I}[\alpha,s]$-skeleton $\hat{\mathbb{H}} \subseteq_w \mathbb{C}\hat{\mathbb{G}}$ violates the freeness condition through some $\hat{g} \in g_1\hat{\mathbb{G}}[\alpha_1] \cap g_2\hat{\mathbb{G}}[\alpha_2] \neq \emptyset$ where $\mathbb{C}\hat{\mathbb{G}}[\mathbb{I},\alpha_1,s_1;g_1] \cap \mathbb{C}\hat{\mathbb{G}}[\mathbb{I},\alpha_2,s_2;g_2] = \emptyset$. By assumption h is injective in restriction to each one of the two cosets $g_1\hat{\mathbb{G}}[\alpha_i]$ and $g_2\hat{\mathbb{G}}[\alpha_2]$. By Lemma 3.6 $\hat{\mathbb{G}}[\alpha]$ is 2-acyclic so that these two cosets intersect precisely in the coset $g\mathbb{G}[\alpha_0]$ for $\alpha_0 = \alpha_1 \cap \alpha_2$. By 2-acyclicity, h must be injective in restriction to the union of these overlapping cosets. Therefore the violation of freeness in $\mathbb{CG}[\alpha]$ would be isomorphically mapped onto a violation in $\mathbb{G}[\alpha]$. ☐

An immediate observation relates freeness to reachability notions that are crucial in connection with the separation condition for coset cycles. To make the connection with Lemma 7.7 below, consider α,α_1,α_2 to play the rôles of $\alpha_i,\alpha_{i,i-1},\alpha_{i,i+1}$ in a potential coset cycle or $\mathbb{I}$-coset cycle.

Observation 7.5. *For $\alpha_1,\alpha_2 \in \Gamma(\alpha)$ and two vertices g_2,g_2 of the same $\mathbb{I}[\alpha,s]$-skeleton $\mathbb{CG}[\mathbb{I},\alpha,s_1;g_1] = \mathbb{CG}[\mathbb{I},\alpha,s_2;g_2] \subseteq_w g_1\mathbb{CG}[\alpha] = g_2\mathbb{CG}[\alpha]$ in $\mathbb{CG}$, freeness implies that if $\mathbb{CG}[\mathbb{I},\alpha_1,s;g_1] \cap \mathbb{CG}[\mathbb{I},\alpha_2,s';g_2] = \emptyset$ (separation as in $\mathbb{I}$-coset cycles) then also $g_1\mathbb{G}[\alpha_1] \cap g_2\mathbb{G}[\alpha_2] = \emptyset$ (separation as in plain coset cycles).*

Finally, the condition on overlaps in Definition 7.3 takes an especially neat and concrete form if $\mathbb{G}$ is 2-acyclic.

Observation 7.6. *Let the embedded $\mathbb{I}[\alpha,s]$-skeleton $\mathbb{CG}[\mathbb{I},\alpha,s;g] \subseteq_w \mathbb{CG}$ be free in $g\mathbb{CG}[\alpha]$ and let $\mathbb{G}[\alpha]$ be 2-acyclic. Then two cosets $g_1\mathbb{G}[\alpha_i]$ and $g_2\mathbb{G}[\alpha_2]$ for $g_i \in \mathbb{CG}[\mathbb{I},\alpha,s;g]$ and $\alpha_i \in \Gamma(\alpha)$ are either disjoint or their overlap is a single coset of the form $g_0\mathbb{G}[\alpha_0]$ for $\alpha_0 = \alpha_1 \cap \alpha_2$ and an element $g_0 \in \mathbb{CG}[\mathbb{I},\alpha,s;g]$ of the skeleton, which is α_i-reachable from g_i in the skeleton.*

Proof. Due to 2-acyclicity a nonempty intersection of α_i-cosets (within $g\mathbb{CG}[\alpha]$) is a single α_0-coset. By freeness it must intersect the embedded $\mathbb{I}[\alpha,s]$-skeleton in an element g_0 that is α_i-reachable from g_i in the skeleton. The converse implication is obvious. ☐

Lemma 7.7. *If $\mathbb{G}$ is N-acyclic (does not admit plain coset cycles of length up to N) and free over $\mathbb{I}$ then $\mathbb{G}$ is Nacyclic over $\mathbb{I}$ (does not admit $\mathbb{I}$-coset cycles of length up to N).*

Proof. It suffices to show that every $\mathbb{I}$-coset cycle in $\mathbb{G}$ induces plain coset cycles if the links $\mathbb{CG}[\mathbb{I}, \alpha_i, s_i; g_i]$ are extended to the surrounding $g_i\mathbb{CG}[\alpha_i]$. The crux is the separation condition: separation w.r.t. $\mathbb{I}$-reachability in links $\mathbb{CG}[\mathbb{I}, \alpha_i, s_i; g_i] \subseteq_w g_i\mathbb{CG}[\alpha_i]$ needs to translate into the stronger separation condition for the encompassing $g_i\mathbb{G}[\alpha_i]$ themselves. This is precisely the content of Observation 7.5. ☐

7.1 Small coset extensions

Towards the construction of E-groups that are free over $\mathbb{I}$ we construct auxiliary E-graphs that reflect the desired freeness condition locally. The rôle of these *small coset extensions* towards freeness over $\mathbb{I}$ is similar to the rôle of free amalgamation chains for N-acyclicity. They reflect the desired freeness condition locally as a result of an unfolding w.r.t. cosets (cf. Definitions 7.8 and 7.10); and they serve as components in E-graphs $\mathbb{H}$ which globally force the required freeness in $\mathbb{G} = \text{sym}(\mathbb{H})$ (cf. Lemma 7.9).

Recall that $\Gamma(\alpha) = \{\alpha' \subseteq E \colon \alpha' \subsetneq \alpha\}$ with corresponding families of generated subgroups $\Gamma(\mathbb{G}[\alpha])$ and Cayley graphs $\Gamma(\mathbb{CG}[\alpha])$ from a given E-group $\mathbb{G}$. Also recall that we say that $\Gamma(\mathbb{G}[\alpha])$ is compatible with an E-graph if every member of this family, i.e. every $\mathbb{G}[\alpha']$ for $\alpha' \subsetneq \alpha$, is.

Definition 7.8. [small coset extension]
Let $\mathbb{G}$ be an E-group with Cayley graph $\mathbb{CG}$, $\alpha \subseteq E$ and $\mathbb{H} = \mathbb{H}[\mathbb{I}, s, \alpha; v]$ an $\mathbb{I}[\alpha, s]$-skeleton. An α-graph $\hat{\mathbb{H}}$ is a *small coset extension* of $\mathbb{H}$ w.r.t. $\mathbb{G}$ (more specifically: w.r.t. $\Gamma(\mathbb{G}[\alpha])$) if

- (i) $\mathbb{H} \subseteq_w \hat{\mathbb{H}}$,
- (ii) for $\alpha' \in \Gamma(\alpha)$ every α'-connected component of $\mathbb{H}$ is contained in a weak substructure of $\hat{\mathbb{H}}$ that is isomorphic to $\mathbb{CG}[\alpha']$.

The small coset extension $\mathbb{H} \subseteq_w \hat{\mathbb{H}}$ is *free* if

- (iii) for all $\alpha_1, \alpha_2 \in \Gamma(\alpha)$, any two α_i-connected components of $\mathbb{H}$ that are disjoint in $\mathbb{H}$ extend into disjoint α_i-connected components of $\hat{\mathbb{H}}$.

Note that freeness as in (iii) implies that the weak substructure relationship in (i) becomes an induced substructure relationship: an e-edge of $\hat{\mathbb{H}}$ between vertices of $\mathbb{H}$ is an $\{e\}$-connected component of $\hat{\mathbb{H}}$ and, by (iii), must be present in $\mathbb{H}$. The existence of small coset extensions (not necessarily free) is obvious e.g. for embedded $\mathbb{I}[\alpha, s]$-skeletons as discussed in Example 7.2. Here the small coset extension $\mathbb{CG}[\mathbb{I}, \alpha, s; g] \subseteq_w \mathbb{H}$ is obtained as a weak substructure $\mathbb{H} \subseteq_w \mathbb{CG}[\alpha]$ formed by the union of the α'-cosets of the elements of $\mathbb{CG}[\mathbb{I}, \alpha, s; g]$, for $\alpha' \in \Gamma(\alpha)$. This small coset extension is free if the skeleton is free in $\mathbb{CG}[\alpha]$.

Lemma 7.9. *Let $\mathbb{G}$ be an E-group, $\alpha \subseteq E$ such that $\mathbb{G}[\alpha]$ is N-acyclic for some $N \geqslant 2$. Let $\mathbb{H}$ comprise, as components, a small coset extension of an $\mathbb{I}[\alpha, s]$-skeleton w.r.t. $\mathbb{G}$ that is free and an isomorphic copy of $\mathbb{CG}$. Assume further that $\Gamma(\mathbb{G}[\alpha])$ is compatible with $\mathbb{H}$. Then $\hat{\mathbb{G}} := \text{sym}(\mathbb{H})$ is such that*

(i) $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$, $\hat{\mathbb{G}}[\alpha'] \simeq \mathbb{G}[\alpha']$ for all $\alpha' \in \Gamma(\alpha)$, and $\hat{\mathbb{G}}[\alpha]$ is N-acyclic;

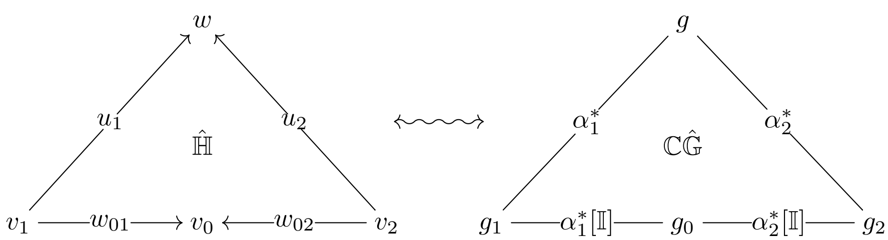

Figure 1: Patterns of walks in union of relevant α_i-components of $\hat{\mathbb{H}}$ and $\mathbb{C}\hat{\mathbb{G}}$.

(ii) the embedded $\mathbb{I}[\alpha, s]$*-skeleton* $\mathbb{C}\hat{\mathbb{G}}[\mathbb{I}, s, \alpha] \subseteq_w \mathbb{C}\hat{\mathbb{G}}[\alpha]$ *is free in* $\mathbb{C}\hat{\mathbb{G}}[\alpha]$.

Proof. Let $\hat{\mathbb{H}}_0$ be the free small coset extension w.r.t. $\mathbb{G}$ that appears as a component of $\mathbb{H}$, $\mathbb{H}_0 \subseteq_w \hat{\mathbb{H}}_0$ the underlying $\mathbb{I}[\alpha, s]$-skeleton.

In the situation of the lemma $\mathbb{H}$ also has a component $\mathbb{C}\mathbb{G}$ so that, $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$ and, as $\Gamma(\mathbb{G}[\alpha])$ is compatible with $\mathbb{H}$, $\hat{\mathbb{G}}[\alpha'] \simeq \mathbb{G}[\alpha']$ for $\alpha' \in \Gamma(\alpha)$ (cf. Lemma 2.11, or its straightforward analogue for $\Gamma(\mathbb{G}[\alpha])$ instead of $\Gamma_k(\mathbb{G})$). It follows that $\hat{\mathbb{G}}[\alpha]$ inherits N-acyclicity from $\mathbb{G}[\alpha]$ (cf. Lemma 3.6).

It remains to argue for freeness of the embedded $\mathbb{I}[\alpha, s]$-skeleton

$$\mathbb{C}\hat{\mathbb{G}}[\mathbb{I}, \alpha, s] := \mathbb{C}\hat{\mathbb{G}}[\mathbb{I}, \alpha, s; 1] \subseteq_w \mathbb{C}\hat{\mathbb{G}}[\alpha].$$

Let $\alpha_1, \alpha_2 \in \Gamma(\alpha)$ and consider the freeness condition for $g_1\mathbb{C}\hat{\mathbb{G}}[\alpha_1]$ and $g_2\mathbb{C}\hat{\mathbb{G}}[\alpha_2]$ where $g_1, g_2 \in \mathbb{C}\hat{\mathbb{G}}[\mathbb{I}, \alpha, s]$. Compare Figure 1, with g_1, g_2, g on the right-hand side. Let $w_i \in \alpha^*[\mathbb{I}, s]$ be such that $g_i = [w_i]_{\hat{\mathbb{G}}}$ and $u_i \in \alpha_i^*$ such that $g_1[u_1]_{\hat{\mathbb{G}}} = g_2[u_2]_{\hat{\mathbb{G}}} =: g \in g_1\hat{\mathbb{G}}[\alpha_1] \cap g_2\hat{\mathbb{G}}[\alpha_2]$. Considering the operation of $w_i u_i$ on $v \in \mathbb{H}_0 \subseteq_w \hat{\mathbb{H}}_0$, we note that the w_i map v to vertices $v_i \in \mathbb{H}_0$ that are reached from v on corresponding $\alpha^*[\mathbb{I}, s]$-walks. Since $[w_1u_1]_{\hat{\mathbb{G}}} = g = [w_2u_2]_{\hat{\mathbb{G}}}$, the u_i-walks from the vertices v_i in $\hat{\mathbb{H}}_0$ must meet in the same vertex $w \in \hat{\mathbb{H}}$, which therefore lies in the intersection of the α_i-connected components of the v_i in $\hat{\mathbb{H}}_0$. Compare the left-hand side of Figure 1. Since $\hat{\mathbb{H}}_0 \;{}_w\!\supseteq \mathbb{H}_0$ is a free small coset extension, the α_i-connected components of the v_i in $\mathbb{H}_0$ must intersect in some vertex v_0 of $\mathbb{H}_0$. It follows that v_0 is reachable from v_i on an $\alpha_i^*[\mathbb{I}, s_i]$-walk, for $i = 1, 2$; let these walks be labelled by $w_{0i} \in \alpha_i^*[\mathbb{I}, s_i]$ for the appropriate $s_i \in \mathbb{I}[\alpha, s]$. The union of the α_i-connected components of v_i in $\hat{\mathbb{H}}$ contains v_0 and w (both in the intersection of the two components) as well as v_1 in the former and v_2 in the latter of these components. By Observation 4.2 this union is isomorphic, as an E-graph, to the free amalgam $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2] \simeq \mathbb{C}\hat{\mathbb{G}}[\alpha_1] \oplus \mathbb{C}\hat{\mathbb{G}}[\alpha_2]$. It follows that also in $\hat{\mathbb{G}}$ the word

$$w_{01}^{-1} u_1 u_2^{-1} w_{02}$$

labels a cycle. In terms of the natural isomorphisms, this labelling traces out the cycles

$$\begin{aligned}
&(v_0, v_1, w, v_2, v_0)\\
\simeq\ &(1, [w_{01}]_{\hat{\mathbb{G}}}^{-1}, [w_{01}]_{\hat{\mathbb{G}}}^{-1}[u_1]_{\hat{\mathbb{G}}}, [w_{01}]_{\hat{\mathbb{G}}}^{-1}[u_1]_{\hat{\mathbb{G}}}[u_2]_{\hat{\mathbb{G}}}^{-1}, 1)\\
\simeq\ &(g_0, g_1, g, g_2, g_0)
\end{aligned}$$

for $g_0 := g_1[w_{01}]_{\hat{\mathbb{G}}}$. This translates into the equality

$$[w_{01}]_{\hat{\mathbb{G}}}^{-1}[u_1]_{\hat{\mathbb{G}}}[u_2]_{\hat{\mathbb{G}}}^{-1}[w_{02}]_{\hat{\mathbb{G}}} = 1$$

in $\hat{\mathbb{G}}$, which implies that $g_1^{-1}g_2 = [u_1]_{\hat{\mathbb{G}}}[u_2]_{\hat{\mathbb{G}}}^{-1} = [w_{01}]_{\hat{\mathbb{G}}}[w_{02}]_{\hat{\mathbb{G}}}^{-1}$ or that

$$g_0 = g_1[w_{01}]_{\hat{\mathbb{G}}} \in g_1\mathbb{C}\hat{\mathbb{G}}[\mathbb{I}, s_1, \alpha_1] \cap g_2\mathbb{C}\hat{\mathbb{G}}[\mathbb{I}, s_2, \alpha_2] \neq \emptyset$$

as $[w_{0i}]_{\hat{\mathbb{G}}} \in \alpha_i^*[\mathbb{I}, s_i]$. So the intersection of the cosets $g_i\mathbb{C}\hat{\mathbb{G}}[\alpha_i]$ does conform to the condition imposed by freeness according to Definition 7.3. □

We turn to a canonical construction of a free coset extension as an amalgam over the given skeleton. The construction itself, as a natural quotient structure, is comparatively straightforward. The compatibility requirement in Lemma 7.9 will then be the major challenge towards its usefulness.

In connection with the precondition on freeness at the level of subgroups in $\Gamma(\mathbb{G}[\alpha])$ compare Definition 7.3.

Definition 7.10. [small coset amalgam]
Let $\alpha \subseteq E$ and $\mathbb{G}$ be an E-group for which all subgroups in $\Gamma(\mathbb{G}[\alpha])$ are 2-acyclic and free over $\mathbb{I}$. Let $\mathbb{H}$ be an $\mathbb{I}[\alpha, s]$-skeleton such that the α'-connected components of $\mathbb{H}$ are isomorphic to corresponding embedded $\mathbb{I}[\alpha', s']$-skeletons in the E-group $\mathbb{G}$ for all $\alpha' \in \Gamma(\alpha)$. In this situation the *small coset amalgam* $\mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha)$ is defined as the quotient

$$\mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha) := \Big(\textstyle\bigcup_{v\in\mathbb{H}, \alpha'\in\Gamma(\alpha)} (\mathbb{C}\mathbb{G}[\alpha'] \times \{(v, \alpha')\})\Big) / \approx$$

where $\approx$ is the equivalence relation induced by identifications of g_1 in the (v_1, α_1)-tagged copy of $\mathbb{C}\mathbb{G}[\alpha_1]$ and g_2 in the (v_2, α_2)-tagged copy of $\mathbb{C}\mathbb{G}[\alpha_2]$ if there is a vertex in the intersection of the α_i-components of the v_i in $\mathbb{H}$ that forces $g_1 = g_2$ in the overlap of copies of $\mathbb{C}\mathbb{G}[\alpha_i]$ (see detailed discussion below).

Note that the isomorphism type of $\mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha)$ is fully determined by $\mathbb{H}$ and the $\Gamma(\mathbb{C}\mathbb{G}[\alpha])$, i.e. by the $\mathbb{C}\mathbb{G}[\alpha']$ for $\alpha' \subsetneq \alpha$, rather than the global structure of $\mathbb{C}\mathbb{G}$. This is relevant for considerations similar to Lemma 2.11.

By definition $\approx$ is the transitive closure of $\sim$, which is defined as follows (cf. Figure 2):

(†)

$(g_1, v_1, \alpha_1) \sim (g_2, v_2, \alpha_2)$

if there are $w_i \in \alpha_i^*[\mathbb{I}, s_i]$ for suitable s_i, and $w \in (\alpha_1 \cap \alpha_2)^*$ such that

(i) $[w_1]_{\mathbb{H}}(v_1) = [w_2]_{\mathbb{H}}(v_2)$ in $\mathbb{H}$,

(ii) $g_i = [w_i w]_{\mathbb{G}} \in \mathbb{G}[\alpha_i]$, for $i = 1, 2$.

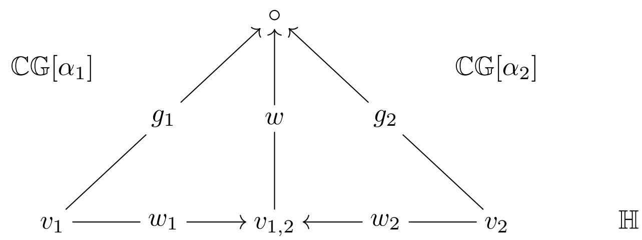

Figure 2: Pattern for identifications in $\mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha)$.

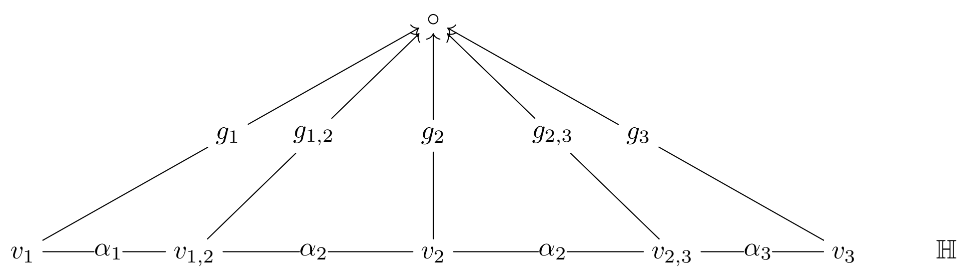

Figure 3: Pattern for two consecutive identifications.

Here (i) says that the w_i provide the addresses of the shared vertex of $\mathbb{H}$ in the copies of the $\mathbb{CG}[\alpha_i]$, relative to their anchor points v_i that get identified with 1 in $\mathbb{CG}[\alpha_i]$. This shared vertex gives rise to a forced identification according to (ii). Conditions (i)–(ii) can be rephrased as the existence of

(a) a vertex $v_{1,2}$ in the intersection of the α_i-components of the v_i in $\mathbb{H}$,

(b) some $g_0 \in \mathbb{CG}[\alpha_0]$ for $\alpha_0 := \alpha_1 \cap \alpha_2$,

such that $(g_0, v_{1,2}, \alpha_0)$ is identified with both (g_i, v_i, α_i) within their copies of $\mathbb{CG}[\alpha_i]$, which are attached at v_i. Compare Figure 2, where we read off $v_{1,2} = [w_1]_{\mathbb{H}}(v_1) = [w_2]_{\mathbb{H}}(v_2)$ and $g_0 = [w]_{\mathbb{G}}$.

The precondition

$$(\ddagger) \qquad \text{the subgroups in } \Gamma(\mathbb{G}[\alpha]) \text{ are free over } \mathbb{I}$$

is essential for the analysis of $\approx$. As $\approx$ is the transitive closure of $\sim$ according to (†), we examine the effect successive identifications through $\sim$. For the following compare Figure 3. Let $(g_1, v_1, \alpha_1) \sim (g_1, v_1, \alpha_1)$, and $(g_2, v_2, \alpha_2) \sim (g_3, v_3, \alpha_3)$ according to (i)–(ii) in (†), for $\alpha_i \in \Gamma(\alpha)$, i.e. $\alpha_i \subsetneq \alpha$. Using the rephrasing according to (a)–(b) above, this implies the existence of intermediaries $v)_{i,j}$ and $g_{i,j}$ for $(i,j) = (1,2)$ and $(i,j) = (2,3)$ where

(a) $v_{i,j}$ is in the intersection of the α_k-components of v_k for $k = i, j$,

(b) $g_{i,j} \in \mathbb{G}[\alpha_{i,j}]$ with $\alpha_{i,j} := \alpha_i \cap \alpha_j$

such that (cf. Figure 3)

– $(v_{1,2}, g_{1,2}, \alpha_{1,2})$ is identified with (v_1, g_1, α_1) (in the $\mathbb{CG}[\alpha_1]$-copy at v_1) and with (v_2, g_2, α_2) (in the $\mathbb{CG}[\alpha_2]$-copy at v_2);
– $(v_{2,3}, g_{2,3}, \alpha_{2,3})$ is identified with (v_2, g_2, α_2) (in the $\mathbb{CG}[\alpha_2]$-copy at v_2) and with (v_3, g_3, α_3) (in the $\mathbb{CG}[\alpha_3]$-copy at v_3).

We look at those identifications that occur in the central $\mathbb{CG}[\alpha_2]$-copy. That the $\alpha_2^*[\mathbb{I}]$-component of $\mathbb{H}$ containing $v_{1,2}, v_2$ and $v_{2,3}$ is free in this $\mathbb{CG}[\alpha_2]$-copy according to (‡) means that at least one of the following must hold

(1) not $\alpha_{1,2} \subsetneq \alpha_2$, i.e. $\alpha_2 \subseteq \alpha_1$, or
(2) not $\alpha_{2,3} \subsetneq \alpha_2$, i.e. $\alpha_2 \subseteq \alpha_3$, or
(3) the $\alpha_{1,2}$-component of $v_{1,2}$ and the $\alpha_{2,3}$-component of $v_{2,3}$ intersect in the α_2-component of v_2 in $\mathbb{H}$.

In either one of the first two cases the identification of (v_1, g_1, α_1) with (v_3, g_3, α_3) is obviously mediated by a direct identification through $\sim$ as defined in (†). The same is true in the third case, since it implies that also the α_1-component of v_1 and the α_2-component of v_3 in $\mathbb{H}$ intersect in some vertex v_2' for which there is a $g_2' \in \mathbb{G}[\alpha_{1,2} \cap \alpha_{2,3}]$ that mediates a direct identification according to (a)–(b) above. So the crucial assumption (‡) implies that $\sim$ is itself transitive, hence equal to $\approx$.

Parts (i) and (ii) in the following are an analogue of the assertion in Remark 4.7 for small coset amalgams instead of amalgamation clusters.

Remark 7.11. *For $\hat{\mathbb{H}} := \mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha) \;{}_w\!\supseteq \mathbb{H}$ as in Definition 7.10 the relation $\sim$ in (†) is itself an equivalence relation so that $\hat{\mathbb{H}}$ is the quotient w.r.t. $\sim$. Moreover, for every element x of $\hat{\mathbb{H}}$ there is*

(i) a unique minimal set $\alpha_x \in \Gamma(\alpha)$ for which x is represented in a copy of $\mathbb{CG}[\alpha_x]$ attached to some element $v \in \mathbb{H}$, and
(ii) a single full α_x-component of $\mathbb{H}$ consisting of those elements $v \in \mathbb{H}$ that admit a representation of x by (g, v, α_x) for some $g \in \mathbb{G}[\alpha_x]$.

Proof. For (i) note that α_x can be obtained as the intersection of all generator sets $\alpha' \in \Gamma(\alpha)$ for which the α'-connected component of x in $\hat{\mathbb{H}}$ intersects $\mathbb{H}$. This set is closed under intersections as indicated in the characterisation of $\sim$ in (a)–(b). Tuples that represent x can only be of the from (g, v, α') for $\alpha_x \subseteq \alpha' \in \Gamma(\alpha)$. For (ii), the relevant subset of $\mathbb{H}$ is closed under α_x-reachability in $\mathbb{H}$. That it cannot comprise disjoint α_x-components follows again by definition of $\sim$ in (a)–(b) above, (†) and (‡). □

Lemma 7.12. *In the situation of Definition 7.10 the small coset amalgam $\hat{\mathbb{H}} := \mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha) \;{}_w\!\supseteq \mathbb{H}$ is a free small coset extension of the $\mathbb{I}[\alpha, s]$-skeleton $\mathbb{H}$.*

Proof. Conditions conditions (i)–(iii) on free small coset extensions in Definition 7.8 are directly built into the definition of $\mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha)$. We need to check that it actually is an E-graph.

Consider elements x_i represented by (g_i, v_i, α_i) where $\alpha_i = \alpha_{x_i}$, $g_i \in \mathbb{G}[\alpha_i]$ for $i = 1, 2$ and let (x_1, x_2) be an e-edge in $\hat{\mathbb{H}}$. Let this e-edge stem from a copy of some $\mathbb{CG}[\alpha_0]$ at $v_0 \in \mathbb{H}$, where $e \in \alpha_0$ (all $\alpha_i \in \Gamma(\alpha)$, i.e. $\alpha_i \subsetneq \alpha$). As $x_i \sim (g_{0i}, v_0, \alpha_0)$ for $i = 1, 2$ and suitable $g_{0i} \in \mathbb{G}[\alpha_0]$, we find that $\alpha_i \subseteq \alpha_0$ for $i = 1, 2$, and all v_i are in the same α_0-component of $\mathbb{H}$. We may assume w.l.o.g. that

$$\alpha_0 = \alpha_1 \cup \alpha_2 \cup \{e\},$$

since $\alpha_0 \supseteq \{e\}$ and since this α_0 is sufficient to put x_1, x_2, v_1, v_2 in the same α_0-component of $\hat{\mathbb{H}}$. Looking at the triangle formed by v_0, x_1, x_2 in the $\mathbb{CG}[\alpha_0]$-copy that represents the x_i and the e-link between them, 2-acyclicity of $\mathbb{G}[\alpha_0]$ implies that $e \in \alpha_1 \cup \alpha_2$, and hence that

$$\alpha_0 = \alpha_{x_1} \cup \alpha_{x_2}.$$

We further distinguish two cases, based on whether $e \in \alpha_1 \cap \alpha_2$ or not. If $e \in \alpha_1 \cap \alpha_2$, then

$$\alpha_0 = \alpha_{x_1} = \alpha_{x_2},$$

as $e \in \alpha_1 \cap \alpha_2 = \alpha_{x_1} \cap \alpha_{x_2}$ implies $\alpha_{x_1} = \alpha_{x_2}$ by minimality of these generator sets for x_1 and x_2. Otherwise, i.e. if $e \in \alpha_1 \setminus \alpha_2 = \alpha_{x_1} \setminus \alpha_{x_2}$ or $e \in \alpha_2 \setminus \alpha_1 = \alpha_{x_2} \setminus \alpha_{x_1}$, it must be that

$$\alpha_0 = \alpha_{x_1} = \alpha_{x_2} \mathbin{\dot{\cup}} \{e\}$$
$$\text{or} \quad \alpha_0 = \alpha_{x_2} = \alpha_{x_1} \mathbin{\dot{\cup}} \{e\} :$$

if, for instance, $e \in \alpha_1 \setminus \alpha_2 = \alpha_{x_1} \setminus \alpha_{x_2}$, since $\alpha_2 \cup \{e\} \subseteq \alpha_0 \subsetneq \alpha$, the representations of x_2 relative to v_1 and v_2 show that $\alpha_2 = \alpha_{x_2} \subseteq \alpha_1 \cap (\alpha_2 \cup \{e\})$, whence $\alpha_2 \subseteq \alpha_1 = \alpha_2 \cup \{e\}$.

This means that e-edges can either occur between elements that share the same $\alpha_x \subsetneq \alpha$ with $e \in \alpha_x$, or else between elements x with $e \notin \alpha_x$ and y with $\alpha_y = \alpha_x \cup \{e\}$. In the first case, both vertices belong to the same $\mathbb{CG}[\alpha_x]$-copy in $\hat{\mathbb{H}}$ that comprises this e-edge; in the second case, the e-edge is simultaneously part of any $\mathbb{CG}[\alpha']$-copy in $\hat{\mathbb{H}}$ that contains y, as well as of any $\mathbb{CG}[\alpha']$-copy in $\hat{\mathbb{H}}$ that contains x and has $e \in \alpha'$. It follows in both cases that no further e-edge can be incident with either vertex so that $\hat{\mathbb{H}}$ is indeed an E-graph. So $\hat{\mathbb{H}} = \mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha)$ is a free small coset extension of $\mathbb{H}$. □

Remark 7.13. *In the situation of Definition 7.10 any free small coset extension of* $\mathbb{H}$ *is isomorphic to* $\mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha)$.

Remark 7.14. *If the embedded* $\mathbb{I}[\alpha, s]$*-skeleton* $\mathbb{H} = \mathbb{CG}[\mathbb{I}, \alpha, s] \subseteq_w \mathbb{CG}[\alpha]$ *is free in* $\mathbb{CG}[\alpha]$*, then* $\mathbb{CG}[\alpha]$ *contains, as a weak subgraph, a free small coset extension* $\hat{\mathbb{H}} \;{}_w\!\supseteq \mathbb{H}$. *If* $\mathbb{G}[\alpha]$ *is 2-acyclic, then this* $\hat{\mathbb{H}}$ *is isomorphic to* $\mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha)$ *and we refer to this realisation of the free small coset extension* $\mathsf{CE}(\mathbb{CG}[\mathbb{I}, \alpha, s], \mathbb{G}, \alpha) \subseteq_w \mathbb{CG}[\alpha]$ *as the* embedded *free small coset extension.*

7.2 Cluster property for small coset amalgams

The next goal is to establish that, under suitable conditions on $\Gamma(\mathbb{G}[\alpha])$, the free coset extension $\mathsf{CE}(\mathbb{H},\mathbb{G},\alpha)$ also satisfies the compatibility conditions that are relevant for Lemma 7.9. For compatibility of $\Gamma(\mathbb{G}[\alpha])$ with $\mathsf{CE}(\mathbb{H},\mathbb{G},\alpha)$ we look at suitable weak substructures that enclose β-connected components for $\beta \in \Gamma(\alpha)$. It turns out that free amalgamation clusters (cf. Definition 4.5) are suitable for this purpose. Assuming corresponding behaviour at the level of $\Gamma(\mathbb{G}[\alpha])$, every β-component of $\mathsf{CE}(\mathbb{H},\mathbb{G},\alpha)$ is contained in a weak substructure that is isomorphic to a free amalgamation cluster of $\Gamma(\mathbb{G}[\alpha])$-cosets. If the subgroups in $\Gamma(\mathbb{G}[\alpha])$ satsify the cluster property of Definition 4.11 then Lemma 4.12 applies to show that this β-component is itself isomorphic to a free amalgamation cluster. The following is an analogue of Definition 4.11.

Definition 7.15. [cluster property for small coset amalgams]
In the situation of Definition 7.10 let $\hat{\mathbb{H}} = \mathsf{CE}(\mathbb{H},\mathbb{G},\alpha)$ be the small coset extension of the $\mathbb{I}[\alpha,s]$-skeleton $\mathbb{H}$. We say that this small coset extension $\hat{\mathbb{H}} \;{}_w{\supseteq}\; \mathbb{H}$ has the *cluster property* if every β-connected component B of $\hat{\mathbb{H}}$ for $\beta \in \Gamma(\alpha)$ satisfies the following:

(i) B has an element x such that $\alpha_x = \alpha_B := \bigcap\{\alpha_x : x \in B\}$ and a constituent $\mathbb{CG}[\alpha_B]$-copy of $\hat{\mathbb{H}}$ containing x that is contained in every constituent $\mathbb{CG}[\alpha']$-copy of $\hat{\mathbb{H}}$ that intersects B;

(ii) B is contained in a weak substructure of $\hat{\mathbb{H}}$ that is isomorphic to a free amalgamation cluster $\bigoplus_{x\in B} \mathbb{CG}[\alpha_x]$ of $\mathbb{CG}[\alpha_x]$ from $\Gamma(\mathbb{CG}[\alpha])$, whose core is the constituent $\mathbb{CG}[\alpha_B]$-copy from (i).

The following lemma shows that the cluster property is preserved in the passage from free small coset extensions at the level of the smaller subgroups in $\Gamma(\mathbb{G}[\alpha])$ to any free small coset extension obtained as a small coset amalgam according to Definition 7.10. It is an analogue of Lemma 4.12 for small coset amalgams. If compatibility of $\Gamma(\hat{\mathbb{G}}[\alpha])$ with free amalgamation chains of length N and free amalgamation clusters is simultaneously maintained in line with Lemmas 4.10 and 4.12, the cluster property implies compatibility as required for Lemma 7.9. It follows that Lemma 7.9 supports an inductive construction of E-groups that are N-acyclic, free and satisfy the cluster property for free small coset amalgams. Compare Remark 7.13 on embedded free small coset extensions.

Lemma 7.16. *In the situation of Definition 7.10 consider the free small coset extension $\hat{\mathbb{H}} = \mathsf{CE}(\mathbb{H},\mathbb{G},\alpha)$. If all the embedded free small coset extensions $\mathbb{CG}[\mathbb{I},\alpha',s'] \subseteq_w \mathsf{CE}(\mathbb{CG}[\mathbb{I},\alpha',s'],\mathbb{G}[\alpha'],\alpha') \subseteq_w \mathbb{CG}[\alpha']$ for embedded $\mathbb{I}[\alpha',s']$-skeletons $\mathbb{CG}[\mathbb{I},\alpha',s']$ and $\alpha' \in \Gamma(\alpha)$ satisfy the cluster property, then $\hat{\mathbb{H}}$ has the cluster property.*

Proof. Let $\beta \subsetneq \alpha$ and consider a β-component B of $\hat{\mathbb{H}}$. If B is contained in some $\mathbb{CG}[\alpha']$-copy of $\hat{\mathbb{H}}$ for $\alpha' \in \Gamma(\alpha)$, we are done: the above analysis of e-edges of $\hat{\mathbb{H}} = \mathsf{CE}(\mathbb{H},\mathbb{G},\alpha)$ shows that every β-edge between vertices of B must be present in $\mathbb{CG}[\alpha']$. So B is a $(\beta \cap \alpha')$-component of a $\mathbb{CG}[\alpha']$-copy and has the required structure by assumption. (This comprises all instances in which B intersects the skeleton $\mathbb{H}$, as then B corresponds to a constituent $\mathbb{CG}[\beta]$-copy of $\hat{\mathbb{H}}$.)

Otherwise consider an element $x \in B$ that is an element of at least two distinct constituent $\mathbb{CG}[\alpha_i]$-copies of $\hat{\mathbb{H}}$, $\alpha_i \in \Gamma(\alpha)$ for $i = 1, 2$, neither of which covers all of B. Consider the set of generators $\alpha_x \subseteq \alpha_0 := \alpha_1 \cap \alpha_2$ and, for $\beta_i := \beta \cap \alpha_i$, the β_i-connected components B_i of x in the constituent $\mathbb{CG}[\alpha_i]$-copies. Clearly $B_i \subseteq B$ and by assumption on $\mathbb{G}[\alpha_i]$, B_i is contained in a copy of a free amalgamation cluster centred on a core $\mathbb{CG}[\alpha_{B_i}] \subseteq_w \mathbb{CG}[\alpha_i]$ where $\alpha_{B_i} = \bigcap_{y \in B_i} \alpha_y$. Let $x_i \in B_i$ be such that $\alpha_{x_i} = \alpha_{B_i}$. By minimality of α_{B_i}, $\alpha_{x_i} = \alpha_{B_i} \subseteq \alpha_x$. Now x_i is linked to x by a β_i-walk, and also (via 1 in the copy of $\mathbb{CG}[\alpha_i]$) by the composition of an α_{x_i}-walk and an α_x-walk, which is an α_0-walk (as $\alpha_{x_i} = \alpha_{B_i} \subseteq \alpha_x \subseteq \alpha_0$). Since $\mathbb{G}[\alpha_i]$ is 2-acyclic, x_i must be linked to x by a β_0-walk for $\beta_0 := \beta \cap \alpha_0$. It follows that both x_i are elements of the shared copy of $\mathbb{CG}[\alpha_0]$. So the core $\mathbb{CG}[\alpha_{B_i}]$-copies for the B_i touch the same β-component of x also in this shared copy of $\mathbb{CG}[\alpha_0]$. For this β-component B_0 of x in $\mathbb{CG}[\alpha_0]$ the assumption on $\mathbb{CG}[\alpha_0]$ gives us the minimal $\alpha_{B_0} = \bigcap_{y \in B_0} \alpha_y$. Arguing with minimality of α_{B_0} in the shared copy of $\mathbb{CG}[\alpha_0]$, which represents x and both x_i, we find that $\alpha_{B_0} \subseteq \alpha_{B_1} \cap \alpha_{B_2}$. By minimality of the α_{B_i} in their copies of $\mathbb{CG}[\alpha_i]$ it follows that $\alpha_{B_1} = \alpha_{B_0} = \alpha_{B_2}$. So B_1 and B_2 share a β_0-coset in a $\mathbb{CG}[\alpha_{B_0}]$-copy of $\hat{\mathbb{H}}$. The generator set α_{B_0} and the constituent $\mathbb{CG}[\alpha_{B_0}]$ of $\hat{\mathbb{H}}$ are uniquely determined by B: $\alpha_{B_1} = \alpha_{B_0} = \alpha_{B_2} = \alpha_{B_i}$ for any α_i arising from the corresponding $B_i \subseteq B$ that contribute to B, by iteration of the above argument. It follows that the same $\mathbb{CG}[\alpha_{B_0}]$-copy and β_0-coset are contained in all $B' \subseteq B$ that arise as β-components of elements of B in their respective $\mathbb{CG}[\alpha']$-copies. Hence B is contained in a free amalgamation cluster of constituents from $\Gamma(\mathbb{CG}[\alpha])$ with this $\mathbb{CG}[\alpha_{B_0}]$-copy as its core. □

Corollary 7.17. *In the situation of Definition 7.10, assume that the subgroups in* $\Gamma(\mathbb{G}[\alpha])$ *are free over* $\mathbb{I}$ *with embedded free small coset extensions that satisfy the cluster property as in Lemma 7.16 and are compatible with free amalgamation clusters of Cayley graphs from* $\Gamma(\mathbb{CG}[\alpha])$*. Then the family* $\Gamma(\mathbb{G}[\alpha])$ *is compatible with the free small coset extension* $\hat{\mathbb{H}} = \mathsf{CE}(\mathbb{H}, \mathbb{G}, \alpha) \;_w\!\supseteq \mathbb{H}$.

Proof. Given that β-connected components of $\hat{\mathbb{H}}$ for $\beta \subsetneq \alpha$ are isomorphic to weak subgraphs of free amalgamation clusters from $\Gamma(\mathbb{CG}[\alpha])$ by Lemma 7.16, the claim of the corollary follows with the compatibility result on free amalgamation clusters in Lemma 4.12. □

The following lemma outlines an inductive construction of E-groups that are N-acyclic over $\mathbb{I}$; it is an analogue of Lemma 5.2, which provided a similar construction for plain coset acyclicity. That is now also comprised as a special case, viz. for the trivial constraint graph with just loops on a single vertex. Recall that Γ_k stands for the set of generator subsets $\alpha \subseteq E$ of size $|\alpha| < k$, and $\Gamma_k(\mathbb{G})$ for the family of generated subgroups $\mathbb{G}[\alpha]$ generated by $\alpha \in \Gamma_k$. The construction of finite E-groups $\mathbb{G}$ that are N-acyclic over $\mathbb{I}$ is again achieved in an inductive process that produces a sequence of E-groups $\mathbb{G}_k$, for increasing values of k in the range between 1 and $|E|$ such that

$$\mathbb{G}_k \text{ is } N\text{-acyclic w.r.t. } \mathbb{I}[\alpha]\text{-cosets for } \alpha \in \Gamma_k.$$

These levels of N-acyclicity over $\mathbb{I}$ arise as a consequence of the more specific sufficient condition that

$$\mathbb{G}_k \text{ is free over } \mathbb{I} \text{ and } N\text{-acyclic w.r.t. } \Gamma_k(\mathbb{G}_k).$$

As before, $\mathbb{G}_{k+1}$ is obtained as $\mathbb{G}_{k+1} := \mathrm{sym}(\mathbb{H}_k)$ from an E-graph $\mathbb{H}_k$ defined from $\mathbb{C}\mathbb{G}_k$. And, just as in the case of Lemma 5.2, $\mathbb{H}_k$ is chosen such that

– $\mathbb{G}_{k+1} \succcurlyeq \mathbb{G}_k$,
– $\Gamma_{k+1}(\mathbb{G}_{k+1}) = \Gamma_{k+1}(\mathbb{G}_k)$, i.e. $\mathbb{G}_{k+1}[\alpha] \simeq \mathbb{G}_k[\alpha]$ for $|\alpha| \leqslant k$.

The second condition requires guarantees of compatibility with $\mathbb{H}_k$. This in turn requires downward compatibility for α-connected components of $\mathbb{H}_k$, where $\alpha \in \Gamma_{k+1}$. That is where the analysis of free small coset extensions and free amalgamation clusters become a crucial ingredient. The analysis of free amalgamation chains is already familiar from Section 5, including their use towards plain coset N-acyclicity.

Lemma 7.18. *For any finite E-group $\mathbb{G}$ and $n \geqslant 1$ there is a sequence of finite E-groups $(\mathbb{G}_k)_{k \leqslant |E|}$ starting with $\mathbb{G}_0 := \mathbb{G}$ such that for $k < |E|$:*

(i) $\mathbb{G}_k \preccurlyeq \mathbb{G}_{k+1}$,
(ii) $\Gamma_{k+1}(\mathbb{G}_k) = \Gamma_{k+1}(\mathbb{G}_{k+1})$,
(iii) $\mathbb{G}_{k+1}$ *is compatible with all free amalgamation chains of length up to n and with all free amalgamation clusters over* $\Gamma(\mathbb{G}_{k+1}) = \Gamma_{k+1}(\mathbb{G}_k)$,
(iv) $\mathbb{G}_{k+1}$ *is free over $\mathbb{I}$ and N-acyclic w.r.t.* $\Gamma_{k+1}(\mathbb{G}_{k+1}) = \Gamma_{k+1}(\mathbb{G}_k)$ *for* $N = n + 2$.

It follows that, for $N = n + 2$, each $\mathbb{G}_k[\alpha]$ is N-acyclic over $\mathbb{I}$ for all α of size $|\alpha| \leqslant k$ and that, for $\hat{\mathbb{G}} := \mathbb{G}_{|E|}$,

$$\mathbb{G} \preccurlyeq \hat{\mathbb{G}} \ \ \textit{where } \hat{\mathbb{G}} \textit{ is } N\textit{-acyclic over } \mathbb{I}.$$

Note that free amalgamation clusters include all free amalgamation chains of length 1 and 2 so that the compatibility with chains in (iii) only comes into effect for $n > 2$ (even 4-acyclicity is guranteed without direct recourse to chains, by Lemma 5.1).

Proof. Conditions (ii)–(iv) are trivially satisfied at the level of $k = 0$; condition (i) will, for all k, be guaranteed by letting $\mathbb{G}_{k+1} := \mathrm{sym}(\mathbb{H}_k)$ where $\mathbb{H}_k$ contains as one of its components a copy of $\mathbb{C}\mathbb{G}_k$ (cf. Lemma 2.11). It remains to choose the remaining components of $\mathbb{H}_k$ in such a way that condition (ii) is maintained while (iii) and (iv) are enforced in the inductive generation of the $\mathbb{G}_{k+1}$ for $k \geqslant 0$. In order to implement the first part of (iii) and the plain N-acyclicity requirement of (iv) we let $\mathbb{H}_k$ comprise copies of

– every free amalgamation chain of length up to n
with constituents $\mathbb{G}_k[\alpha_i]$ for $\alpha_i \in \Gamma_{k+1}$,

which is as in Lemma 5.2. To guarantee the second part in (iii) and the freeness requirement in (iv), we now also put, as components of $\mathbb{H}_k$, copies of

– the free small coset amalgam $\mathsf{CE}(\mathbb{CG}_k[\mathbb{I}, \alpha, s], \mathbb{G}_k, \alpha)$
for every $\alpha \in \Gamma_{k+1}$ and $s \in S$,

which is a free small coset extension of the underlying $\mathbb{I}$-skeleton. These ingredients in $\mathbb{H}_k$ are compatible with condition (ii). For the chains this is due to Lemma 4.10 as before. For the free coset extensions it is due to the cluster property, which entails that compatibility by Corollary 7.17.

That plain N-acyclicity is maintained follows from Lemma 5.1; that freeness is maintained follows from Lemma 7.9. □

Similar to Theorem 5.3 we sum up the construction and emphasise its preservation of symmetries over E.

Theorem 7.19. *For every finite* E*-group* $\mathbb{G}$ *that is compatible with the constraint graph* $\mathbb{I}$ *and every* $N \geqslant 2$ *there is a finite* E*-group* $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$ *that is* N*-acyclic over* $\mathbb{I}$ *and fully symmetric over* $\mathbb{G}$ *in the sense that every permutation* $\rho \in \mathrm{Sym}(E)$ *of the generator set* E *that is a symmetry of* $\mathbb{I}$ *and* $\mathbb{G}$ *extends to a symmetry of* $\hat{\mathbb{G}}$*:* $\mathbb{G}^\rho \simeq \mathbb{G} \Rightarrow \hat{\mathbb{G}}^\rho \simeq \hat{\mathbb{G}}$.

In particular, we may obtain finite E-groups $\hat{\mathbb{G}}$ that are N-acyclic over $\mathbb{I}$ and fully symmetric over $\mathbb{I}$ in the sense of admitting every symmetry of $\mathbb{I}$ as a symmetry. For this we may start e.g. from $\mathbb{G} := \mathrm{sym}(\mathbb{H})$ where $\mathbb{H}$ is the disjoint union of $\mathbb{I}$ and the hypercube $\mathbb{H} = 2^E$.

8 From groups to groupoids

In terms of the combinatorial action of the generators $e \in E$ on an E-graph $\mathbb{H}$, and by extension of the monoid structure of E^* on $\mathbb{H}$, the involutive nature of $\pi_e \in \mathrm{Sym}(V)$ is closely tied to the undirected nature of e-edges in E-graphs. We want to overcome this restriction by allowing for directed e-edges. At the same time we may want to relax the strictly prescribed uniformity between vertices. The latter has already been achieved in the context of involutive generators with constraint graphs $\mathbb{I}$ in Section 6. So now we want to allow for vertices of different sorts with *directed* transitions via e-edges between vertices of specific sorts. Some applications of related notions of acyclicity in graph and hypergraph structures inspired by Cayley graphs in [13, 15] are very naturally cast in terms of multi-sorted multi-graph structures and related groupoids. In this section we directly reduce the construction of groupoids with the desired coset acyclicity properties to the constructions of groups from the previous sections.

8.1 Constraint patterns for groupoids

In the following we consider groupoid structures with a specified pattern of sorts (types of elements, objects) and generators (for the groupoidal operation, morphisms). Groupoids in our sense can also be associated with inverse semigroups of correspondingly restricted patterns. We choose a format for the specification of their sorts that is very similar to the format of E-graphs, and call this specification a *constraint pattern*. The corresponding structures generalise the constraint graphs of Section 6 in the desired direction. Such a template will be a directed multi-graph with edge set E and vertex set S, but unlike E-graphs considered so far, the edges $e \in E$ are directed, with an explicit operation of edge reversal.

Definition 8.1. [constraint pattern $\mathbb{I}$]
A *constraint pattern* is a multi-graph $\mathbb{I} = (S, E, \iota_1, \iota_2, \cdot^{-1})$, which we formalise as a two-sorted structure with a set S of vertices and a set E of edges as sorts, linked by surjective maps $\iota_i \colon E \to S$ that associate a source and target vertex with every edge $e \in E$, and a fixpoint-free and involutive operation of edge reversal $e \mapsto e^{-1}$ on E that is compatible with the ι_i in the sense that $\iota_1(e^{-1}) = \iota_2(e)$.

In the following we mostly abbreviate the notation for a constraint pattern $\mathbb{I}$ as above to just $\mathbb{I} = (S, E)$, leaving the remaining structural details implicit.

For $s, s' \in S$, we let $E[s, s'] := \{e \in E \colon \iota_1(e) = s, \iota_2(e) = s'\}$ be the set of edges linking source s to target s'.

In order to extend the notion of $\mathbb{I}$-reachability (based on an undirected constraint graph $\mathbb{I}$ in Section 6) to a similar concept of $\mathbb{I}$-reachability w.r.t. a constraint pattern $\mathbb{I}$ we consider words that label *directed* walks in $\mathbb{I}$. It is constructive to compare related notions in Section 6. All relevant subsets $\alpha \subseteq E$ will in the following always be closed under edge reversal: $\alpha = \alpha^{-1}$. A *reduced* word over E now is a word in which no $e \in E$ is directly followed or preceded by its inverse e^{-1}.

Definition 8.2. [$\mathbb{I}$-words and $\mathbb{I}$-walks]
Natural sets of (reduced) α-words that occur as edge label sequences along directed walks on the constraint pattern $\mathbb{I}$ are defined as follows:

- $\alpha^*[\mathbb{I}] \subseteq E^*$ consists of those $w \in \alpha^*$ that label a walk in $\mathbb{I}$;
- $\alpha^*[\mathbb{I}, s] \subseteq E^*$ consists of those $w \in \alpha^*$ that label a walk from s in $\mathbb{I}$;
- $\alpha^*[\mathbb{I}, s, t] \subseteq E^*$ consists of those $w \in \alpha^*$ that label a walk from s to t in $\mathbb{I}$.

In particular we write $E^*[\mathbb{I}]$ for the set of all (reduced) words over E that label walks in $\mathbb{I}$, and naturally extend the ι-maps to all of $E^*[\mathbb{I}]$ as follows. Since a walk from s to t in $\mathbb{I}$ is a sequence $s = s_0, e_1, s_1, \ldots, e_n, s_n = t$ such that $\iota_1(e_i) = s_{i-1}$ and $\iota_2(e_i) = s_i$ for $1 \leqslant i \leqslant n$, this walk is fully determined by the sequence of edges and can be identified with the word $w = e_1 \ldots e_n \in E^*$. So we think of $E^*[\mathbb{I}]$ as the set of all words $w = e_1 \ldots e_n$ with $\iota_2(e_i) = \iota_1(e_{i+1})$ for $1 \leqslant i < n$, and put $\iota_1(w) := \iota_1(e_1)$ and $\iota_2(w) := \iota_2(e_n)$. This word w labels a walk in $\mathbb{I}$ from the source vertex $\iota_1(w)$ to the target vertex $\iota_2(w)$.

Correspondingly

$$E^*[\mathbb{I}, s, t] = \{w \in E^*[\mathbb{I}] \colon \iota_1(w) = s, \iota_2(w) = t\}.$$

Concatenation between (reduced) words w_1 and w_2 is defined as $w_1 w_2 \in E^*[\mathbb{I}, \iota_1(w_1), \iota_2(w_2)]$ whenever their ι-values match in the sense that $\iota_2(w_1) = \iota_1(w_2)$. This concatenation operation reflects composition of walks in $\mathbb{I}$.

We think of the vertex set S of $\mathbb{I}$ as a set of *sites* or *vertex colours* and of the edge set E as a set of *links* or *edge colours* that will govern the rôles of elements and generators in corresponding groupoids, as in the following definition. A groupoid is viewed as a group-like structure with groupoid elements of specified sorts. These sorts are pairs of sites and specify the source and the target site of the groupoid element. The groupoidal composition operation, which is partial overall, is fully defined for pairs of elements that share the same interface site.

8.2 I-groupoids and their Cayley graphs

Definition 8.3. [$\mathbb{I}$-groupoid]
An $\mathbb{I}$-*groupoid* based on the constraint pattern $\mathbb{I} = (S, E)$ is a groupoid structure of the form $\mathbb{G} = (G, (G_{s,t})_{s,t \in S}, \cdot, (1_s)_{s \in S}, (g_e)_{e \in E})$ where

(i) the family $(G_{s,t})_{s,t \in S}$ partitions the universe G of groupoid elements;[6]

(ii) $\cdot$ is a groupoidal composition operation mapping any pair of elements in $G_{s,t} \times G_{t,u}$ to an element of $G_{s,u}$, for all combinations of $s, t, u \in S$;

(iii) $1_s \in G_{s,s}$ is a left and right neutral element w.r.t. $\cdot$, for every $s \in S$;

(iv) G is generated by the family of pairwise distinct elements $g_e \in G_{\iota_1(e), \iota_2(e)}$ for $e \in E$, where $g_{e^{-1}}$ is the groupoidal inverse of g_e w.r.t. $\cdot$: $g_{e^{-1}} = g_e^{-1}$ in the sense that $g_e \cdot g_{e^{-1}} = 1_s$ for $s = \iota_1(e)$ and $g_{e^{-1}} \cdot g_e = 1_{s'}$ for $s' = \iota_2(e)$.

The set $E^*[\mathbb{I}, s, t]$ of those (reduced) words over E that label walks from s to t in $\mathbb{I}$ now suggests an interpretation of $w \in E^*[\mathbb{I}, s, t]$ as a product of generators that represents a groupoid element in $G_{s,t}$. With $w = e_1 \cdots e_n \in E^*[\mathbb{I}, s, t]$ we associate the groupoid element $[w]_{\mathbb{G}}$ that is the groupoidal composition

$$[w]_{\mathbb{G}} := \prod_{i=1}^n g_{e_i} = g_{e_1} \cdots g_{e_n} \in G_{s,t}.$$

Note that $[w]_{\mathbb{G}} \in G_{s,t}$ precisely for $s = \iota_1(w)$ and $t = \iota_2(w)$. The ι-maps extend to the elements of an $\mathbb{I}$-groupoid $\mathbb{G}$ according to $\iota_i(g) = s_i$ for $i = 1, 2$ if, and only if, $g \in G_{s_1, s_2}$ if, and only if, $g = [w]_{\mathbb{G}}$ for some $w \in E^*[\mathbb{I}, s_1, s_2]$.

There is an obvious notion of homomorphisms between $\mathbb{I}$-groupoids, which needs to respect the rôle of the distinguished generators. The existence of a homomorphism $h \colon \hat{\mathbb{G}} \to \mathbb{G}$ (uniquely determined and surjective if it exists) is expressed as $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$. The following is analogous to Observation 2.3.

Observation 8.4. *The quotient of* $\mathbb{I}$*-walks w.r.t. cancellation of direct edge reversal, as represented by the set of reduced words in* $E^*[\mathbb{I}]$ *as label sequences,*

[6] Some of the sets $G_{s,t}$ may be empty as $\mathbb{I}$ is not required to be connected.

forms an $\mathbb{I}$-groupoid with concatenation. This can be regarded as the free $\mathbb{I}$-groupoid*, which has any other $\mathbb{I}$-groupoid as a homomorphic image.*

Analogous to Definition 2.12 we also consider symmetries that are induced by admissible re-labellings of the sites and links. The natural candidates stem from symmetries of the constraint pattern $\mathbb{I}$. A symmetry of $\mathbb{I}$ is just an automorphisms in the usual sense for $\mathbb{I}$ as a multi-sorted structure, induced by matching permutations of the sets E and S that are compatible with the ι_i and with edge reversal.

Definition 8.5. [symmetry over $\mathbb{I}$]
An automorphism ρ of $\mathbb{I}$ is a *symmetry* of an $\mathbb{I}$-groupoid $\mathbb{G}$ if the renaming of sorts and generators according to ρ, $s \mapsto \rho(s)$ and $g_e \mapsto g_{\rho(e)}$, yields an isomorphic E-groupoid, $\mathbb{G}^\rho \simeq \mathbb{G}$.

In an $\mathbb{I}$-groupoid $\mathbb{G}$, the set $\alpha^*[\mathbb{I}, s, t]$ of (reduced) words over a subset $\alpha = \alpha^{-1} \subseteq E$ carves out a *generated subgroupoid* $\mathbb{G}[\alpha] \subseteq \mathbb{G}$, as well as corresponding *groupoidal cosets* at $g \in G$. These are defined in the obvious manner as

$$\begin{array}{rcl} \mathbb{G}[\alpha] & = & \bigcup_{s,t} \mathbb{G}[\alpha, s, t] \text{ where} \\ \mathbb{G}[\alpha, s, t] & = & \{[w]_{\mathbb{G}} \in G \colon w \in \alpha^*[\mathbb{I}, s, t]\}, \\ \text{and } \ g\mathbb{G}[\alpha] & = & \bigcup_t \{g \cdot [w]_{\mathbb{G}} \colon w \in \alpha^*[\mathbb{I}, \iota_2(g), t]\}. \end{array}$$

As the constraint pattern $\mathbb{I}$ will mostly be fixed, we shall often suppress its explicit mention and write, e.g., just $E^*[s, t]$, or $\alpha^*[s, t]$, just as we already wrote $\mathbb{G}[\alpha]$ or $\mathbb{G}[\alpha, s, t]$ when $\mathbb{I}$ was implicitly determined by $\mathbb{G}$.

The notion of a Cayley graph for a groupoid $\mathbb{G}$ encodes the operation of generators on groupoid elements, by right multiplication, as with Cayley graphs of groups (cf. Definition 2.5).

Definition 8.6. [Cayley graph of an $\mathbb{I}$-groupoid]
The *Cayley graph* of an $\mathbb{I}$-groupoid $\mathbb{G} = (G, (G_{s,t})_{s,t \in S}, \cdot, (1_s)_{s \in S}, (g_e)_{e \in E})$ is the directed edge-coloured graph $\mathbb{C}\mathbb{G} := \mathrm{Cayley}(\mathbb{G}) = (G, (R_e)_{e \in E})$ with vertex set G and edge sets of colour $e \in E$ according to

$$R_e := \{(g, g \cdot g_e) \colon g \in G_{s,t} \text{ for some } s \in S \text{ and } t = \iota_1(e)\}.$$

As with Cayley graphs for E-groups, the Cayley graphs of $\mathbb{I}$-groupoids are more homogeneous than the underlying groupoid. In particular the neutral elements 1_s are not identified in $\mathbb{C}\mathbb{G}$. What is still recognisable in $\mathbb{C}\mathbb{G}$, for an $\mathbb{I}$-groupoid $\mathbb{G}$, is membership in the sets

$$G[s, *] := \iota_1^{-1}(s) \ \text{ and } \ G[*, s] := \iota_2^{-1}(s),$$

which are identified by the existence of corresponding incoming or outgoing R_e-edge for e with $\iota_2(e) = s$ or $\iota_1(e) = s$, respectively. The algebraic structure of the $\mathbb{I}$-groupoid $\mathbb{G}$ is still fully determined by its Cayley graph $\mathbb{C}\mathbb{G}$ in the corresponding action of partial permutations. In the terminology of [11] it can be recovered as a groupoid embedded in the full *symmetric inverse semigroup* $I(G)$ over its vertex set G. In analogy with the case of groups and their Cayley graphs, where the group is realised as a subgroup of the full symmetric group of global permutations of the vertex set the groupoid is realised as a subgroupoid of the set of all bijections between the relevant sets $G[*, s]$.

Observation 8.7. *The $\mathbb{I}$-groupoid $\mathbb{G}$ is isomorphic to the $\mathbb{I}$-groupoid generated by the following bijections π_e for $e \in E[s, s']$:*

$$\begin{array}{rcl} \pi_e \colon G[*, s] & \longrightarrow & G[*, s'] \\ g & \longmapsto & g \cdot g_e, \end{array}$$

where $g \cdot g_e$ is identified as the unique vertex g' of $\mathbb{CG}$ for which $(g, g') \in R_e$. Here $1_s = \mathrm{id}_{G[*,s]}$ is the identity on $G[*, s]$.

Note that π_e is a partial bijection of the set G but total in restriction to the indicated domain and range.

The analogy is carried further in the following (cf. Definitions 2.1, 2.4 and 2.8 in connection with E-graphs and E-groups.) A major difference is the absence of a simple *completion* operation for $\mathbb{I}$-graphs.[7]

Definition 8.8. [$\mathbb{I}$-graph]
An $\mathbb{I}$-*graph*, for a constraint pattern $\mathbb{I} = (S, E)$, is a vertex- and edge-coloured directed graph $\mathbb{H} = (V, (V_s)_{s \in S}, (R_e)_{e \in E})$, whose vertex set V is partitioned into non-empty subsets V_s of vertices of colour $s \in S$, with edge sets $R_e \subseteq V_{\iota_1(e)} \times V_{\iota_2(e)}$ of colour e for $e \in E$ such that $R_{e^{-1}} = R_e^{-1}$. The $\mathbb{I}$-graph $\mathbb{H} = (V, (V_s)_{s \in S}, (R_e)_{e \in E})$ is *complete* if each R_e is a complete matching between $V_{\iota_1(e)}$ and $V_{\iota_2(e)}$ (i.e. the graph of a bijection $\pi_e \colon V_{\iota_1(e)} \to V_{\iota_2(e)}$). An automorphism of $\mathbb{I}$ is a *symmetry* of the $\mathbb{I}$-graph $\mathbb{H}$ if its operation as a renaming on $\mathbb{H}$ yields an isomorphic $\mathbb{I}$-graph: $\mathbb{H}^\rho \simeq \mathbb{H}$.

Clearly the Cayley graph of an $\mathbb{I}$-groupoid is a complete $\mathbb{I}$-graph that shares every symmetry of the $\mathbb{I}$-groupoid (cf. Definition 8.5 for symmetries). Conversely, any complete $\mathbb{I}$-graph determines an $\mathbb{I}$-groupoid in the manner indicated for this special case in Observation 8.7 above. In a complete $\mathbb{I}$-graph $\mathbb{H}$ as in Definition 8.8, the composition of the π_e along $w = e_1 \cdots e_n \in E^*[\mathbb{I}, s, t]$, $\pi_w = \prod_{i=1}^n \pi_{e_i} = \pi_{e_n} \circ \cdots \pi_{e_1}$, induces a bijection $\pi_w \colon V_s \to V_t$, which we denote as $[w]_{\mathbb{H}}$. The natural composition operation on matching interface sites induces the structure of an $\mathbb{I}$-groupoid $\mathbb{G}$ on the set $G = \{\pi_w \colon w \in E^*[\mathbb{I}]\}$.

Definition 8.9. [sym($\mathbb{H}$)]
From a complete $\mathbb{I}$-graph $\mathbb{H}$ as in Definition 8.8, with induced partial bijections π_w for $w \in E^*[\mathbb{I}]$, we obtain the $\mathbb{I}$-groupoid

$$\mathrm{sym}(\mathbb{H}) := \big(G, (G_{s,t}), \cdot\, (1_s), (\pi_e)\big)$$

[7] Indeed the proposal of a naive completion operation accounts for the major flaw in [14].

with $G_{s,t} = \{\pi_w : w \in E^*[\mathbb{I}, s, t]\}$, composition of partial bijections (which is full composition in matching sites to match concatenation of labelling sequences) and identities in corresponding sites as neutral elements.

Observation 8.7 can be restated as $\mathrm{sym}(\mathbb{CG}) \simeq \mathbb{G}$ if $\mathbb{G}$ is an $\mathbb{I}$-groupoid with Cayley graph $\mathbb{CG}$.

Definition 8.10. [compatibility]
An $\mathbb{I}$-groupoid $\mathbb{G}$ is *compatible* with the complete $\mathbb{I}$-graph $\mathbb{H}$ if $\mathbb{G} \succcurlyeq \mathrm{sym}(\mathbb{H})$, i.e. if for all $w \in E^*[\mathbb{I}, s, s]$

$$[w]_{\mathbb{G}} = 1_s \;\Rightarrow\; [w]_{\mathbb{H}} = \mathrm{id}_{V_s}.$$

The following illustrates these concepts and their far-reaching analogy with the situation for E-groups from Section 2.

Observation 8.11. *Any $\mathbb{I}$-groupoid $\mathbb{G}$ is compatible with its Cayley graph. Another $\mathbb{I}$-groupoid $\hat{\mathbb{G}}$ is compatible with the Cayley graph $\mathbb{CG}$ of $\mathbb{G}$ if, and only if, $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$, if, and only if the map $h\colon \hat{\mathbb{G}} \to \mathbb{G}$ which maps $[w]_{\hat{\mathbb{G}}}$ to $[w]_{\mathbb{G}}$ is well-defined (and thus the homomorphism in question).*

8.3 Coset acyclicity for groupoids

Also the following are straightforward analogues of the corresponding notions for E-groups in Definitions 3.1 and 3.2.

Definition 8.12. [coset cycles]
Let $\mathbb{G}$ be an $\mathbb{I}$-groupoid, $n \geqslant 2$. A *coset cycle of length n* in $\mathbb{G}$ is a cyclically indexed sequence of pointed cosets $(g_i\mathbb{G}[\alpha_i], g_i)_{i\in\mathbb{Z}_n}$ such that, for all i,

(i) (connectivity) $g_{i+1} \in g_i\mathbb{G}[\alpha_i]$, i.e. $g_i\mathbb{G}[\alpha_i] = g_{i+1}\mathbb{G}[\alpha_i]$;
(ii) (separation) $g_i\mathbb{G}[\alpha_{i,i-1}] \cap g_{i+1}\mathbb{G}[\alpha_{i,i+1}] = \emptyset$,

where $\alpha_{i,j} := \alpha_i \cap \alpha_j$.

Definition 8.13. [N-acyclicity]
For $N \geqslant 2$, an $\mathbb{I}$-groupoid $\mathbb{G}$ is *N-acyclic* if it admits no coset cycles of lengths up to N.

9 Construction of N-acyclic I-groupoids

We associate with a constraint pattern $\mathbb{I} = (S, E)$ for $\mathbb{I}$-groupoids $\mathbb{G}$ a set $\hat{E}$ of involutive generators and a constraint graph $\hat{\mathbb{I}}$ so that $\mathbb{I}$-groupoids of interest can be identified within suitable $\hat{\mathsf{E}}$-groups $\hat{\mathbb{G}}$ that are compatible with $\hat{\mathbb{I}}$. More specifically, we aim for a low-level interpretation of Cayley graphs of $\mathbb{I}$-groupoids $\mathbb{CG}$ within the direct product of $\hat{\mathbb{I}}$ with the Cayley graph $\mathbb{C}\hat{\mathbb{G}}$ of an $\hat{\mathsf{E}}$-group $\hat{\mathbb{G}}$ that is compatible with $\hat{\mathbb{I}}$. Compare Definition 6.5 for this direct product.

Firstly, we interpret the directed multi-graph structure of the constraint pattern

$$\mathbb{I} = (S, E) = (S, E, \iota_1, \iota_2, \cdot^{-1})$$

in the structure of a constraint graph

$$\hat{\mathbb{I}} = (\hat{S}, (R_{\hat{e}})_{\hat{e} \in \hat{E}})$$

for a set $\hat{E}$ of involutive generators. Recall that the edge relations R_e of the latter are undirected while the edges $e \in E$ of the former are directed. To this end, associate with every $e \in E$ 3 new edge labels $\{e\}$, $\{e, e^{-1}\}$ and $\{e^{-1}\}$ in $\hat{E}$, as well as 2 new vertices s_e and $s_{e^{-1}}$ in $\hat{S}$. On the basis of

$$\begin{array}{rcl} \hat{E} & := & \big\{\{e\}, \{e, e^{-1}\}, \{e^{-1}\} \colon e \in E\big\}, \\ \hat{S} & := & S \cup \big\{s_e, s_{e^{-1}} \colon e \in E\big\}, \end{array}$$

we represent directed e-(multi-)edges as walks of length 3 in an $\hat{\mathsf{E}}$-graph $\hat{\mathbb{I}}$ as follows. We replace the directed edge $e \in E[s, s']$ and its inverse $e' := e^{-1} \in E[s', s]$ in $\mathbb{I}$ by a succession of 3 undirected edges with labels $\{e\}$, $\{e, e'\}$ and $\{e'\}$ that link s and s' via the two new intermediate vertices s_e and $s_{e'}$:

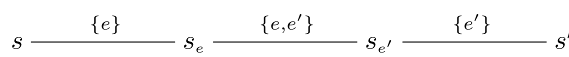

By the same token, a loop $e \in E[s, s]$ at s and its inverse $e' := e^{-1}$ get replaced by a cycle of 3 undirected edges with labels $\{e\}$, $\{e, e'\}$ and $\{e'\}$:

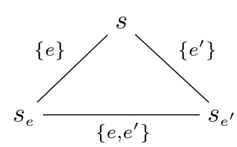

Note that these replacements are inherently symmetric w.r.t. edge reversal in the sense that the replacements really concern the edge pair $\{e, e^{-1}\}$. The direction of e is encoded in the directed nature of the walk

$$s, \{e\}, s_e, \{e, e^{-1}\}, s_{e^{-1}}, \{e^{-1}\}, s',$$

whose reversal exactly is the corresponding walk for e^{-1}. The resulting $\hat{\mathsf{E}}$-graph $\hat{\mathbb{I}}$ is special also in that each one of its edge relations $R_{\hat{e}}$ for $\hat{e} \in \hat{E}$ consists of a single undirected edge. Any automorphism of the constraint pattern $\mathbb{I}$ turns into a symmetry of the $\hat{\mathsf{E}}$-graph $\hat{\mathbb{I}}$, which is the desired constraint graph.

We use this simple schema to associate $\mathbb{I}$-*reachability* w.r.t. the constraint pattern $\mathbb{I}$ for $\mathbb{I}$-groupoids and their Cayley graphs with $\hat{\mathbb{I}}$-*reachability* w.r.t. the constraint graph $\hat{\mathbb{I}}$ for $\hat{\mathsf{E}}$-groups and their Cayley graphs. Overall, this will allow us to directly extract $\mathbb{I}$-groupoids from suitable $\hat{\mathsf{E}}$-groups, in a manner that preserves symmetries and the desired acyclicity properties.

For $\hat{E}, \hat{S}$ and $\hat{\mathbb{I}}$ as just constructed from $\mathbb{I}$, there is a one-to-one correspondence between reduced words in

$$\hat{E}^*[\hat{\mathbb{I}}, s, t] := \{w \in \hat{E}^* \colon\ w \text{ labelling a walk from } s \text{ to } t \text{ in } \hat{\mathbb{I}}\ \}$$

and reduced words in $E^*[\mathbb{I}, s, t]$ that label directed walks from s to t in $\mathbb{I}$. In other words, for all $s, t \in S$, and modulo passage to reduced words, the natural replacement map

$$\begin{array}{rcl} \hat{} : \; E^*[\mathbb{I}, s, t] & \longrightarrow & \hat{E}^*[\hat{\mathbb{I}}, s, t] \\ w = e_1 \cdots e_n & \longmapsto & \hat{w} := \{e_1\}\{e_1, e_1^{-1}\}\{e_1^{-1}\} \cdots \{e_n\}\{e_n, e_n^{-1}\}\{e_n^{-1}\} \end{array}$$

induces a bijection. For this observation it is essential that reduced words in $\hat{E}^*[\hat{\mathbb{I}}]$ can only label walks that link vertices from S if they consist of concatenations of triplets corresponding to admissible orientations of E-edges. In connection with the reduced nature of the words involved, note on one hand that an immediate concatenation of a triplet for $e \in E$ with the triplet for e^{-1} would not be a reduced $\hat{E}$-word. On the other hand, the only non-trivial $\{\{e\}, \{e, e^{-1}\}, \{e^{-1}\}\}$-component of $\hat{\mathbb{I}}$ consists of $\{\iota_1(e), \iota_2(e), s_e, s_{e^{-1}}\}$. The only manner in which a reduced $\hat{E}$-word can label a walk in $\hat{\mathbb{I}}$ that exits this $\{\{e\}, \{e, e^{-1}\}, \{e^{-1}\}\}$-component of $\hat{\mathbb{I}}$ is via $\iota_1(e)$ or $\iota_2(e)$, which are both in S.

For notational convenience we also denote as $\hat{}$ the incarnation of the replacement map at the level of reduced words and at the level of subsets $\alpha \subseteq E$ that are closed under edge reversal:

$$\alpha \;\longmapsto\; \hat{\alpha} := \big\{\{e\}, \{e, e^{-1}\}, \{e^{-1}\} \colon e \in \alpha\big\}.$$

9.1 Groupoids from groups

For a constraint pattern $\mathbb{I} = (S, E)$ and its representation within a constraint graph $\hat{\mathbb{I}}$ for $\hat{\mathsf{E}}$-graphs according to the above translation, consider now an $\hat{\mathsf{E}}$-group $\hat{\mathbb{G}}$ that is compatible with $\hat{\mathbb{I}}$. Let $\mathbb{C}\hat{\mathbb{G}}$ be the Cayley graph of this $\hat{\mathsf{E}}$-group. Recall Definition 6.5 for the definition of a direct product, which we now apply to the $\hat{\mathsf{E}}$-graphs $\hat{\mathbb{I}}$ and $\mathbb{C}\hat{\mathbb{G}}$:

$$\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$$

is an $\hat{\mathsf{E}}$-graph which reflects $\hat{\mathbb{I}}$-reachability in the sense that $(\hat{s}', \hat{g}')$ is in the $\hat{\alpha}$-connected component of $(\hat{s}, \hat{g})$ if, and only if, $\hat{g}'$ is in the $\hat{\mathbb{I}}[\hat{\alpha}, \hat{s}]$-component $\mathbb{C}\hat{\mathbb{G}}[\hat{\mathbb{I}}, \alpha, s; \hat{g}]$ of $\hat{g}$.

We next extract an $\mathbb{I}$-groupoid $\mathbb{G}$ from any $\hat{\mathsf{E}}$-group $\hat{\mathbb{G}}$ that is compatible with the constraint graph $\hat{\mathbb{I}}$. More specifically, the Cayley graph of the target $\mathbb{I}$-groupoid $\mathbb{G} := \hat{\mathbb{G}}[\mathbb{I}]$ is interpreted within the direct product $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$. The idea is to single out the vertices of $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ with $\hat{S}$-component in $S \subseteq \hat{S}$, and to replace $\{e\}\{e, e^{-1}\}\{e^{-1}\}$-walks of length 3 between them by directed E-edges. In essence this is a reversal of the translation that led from E to $\hat{E}$ and from E-graphs to $\hat{\mathsf{E}}$-graphs..

We define $\mathbb{G}$ in terms of its generators $e \in E$, which are interpreted as partial bijections on the vertex set of $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$. We restrict attention to vertices $\{(s, [\hat{w}]_{\hat{\mathbb{G}}}) \in \hat{\mathbb{I}} \otimes \hat{\mathbb{G}} \colon \hat{w} \in \hat{E}^*[\hat{\mathbb{I}}, s, t]\}$ for $s, t \in S \subseteq \hat{S}$, and put

$$G_{s,t} := \{(s, [\hat{w}]_{\hat{\mathbb{G}}}) \colon \hat{w} \in \hat{E}^*[\hat{\mathbb{I}}, s, t]\} = \{(s, [\hat{w}]_{\hat{\mathbb{G}}}) \colon w \in E^*[\mathbb{I}, s, t]\}.$$

The second equality appeals to the identification of reduced words in $\hat{E}^*[\hat{\mathbb{I}}, s, t]$ and $E^*[\mathbb{I}, s, t]$ for $s, t \in S \subseteq \hat{S}$. The sets $G_{s,t}$ are subsets of the vertex set of $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$. They are disjoint by compatibility of $\hat{\mathbb{G}}$ with $\hat{\mathbb{I}}$ and thus partition

$$G := \dot{\bigcup}_{s,t \in S} G_{s,t}$$

into subsets (not all necessarily non-empty unless $\mathbb{I}$ is connected). We write $G_{*,t}$ for the union $G_{*,t} := \bigcup_{s \in S} G_{s,t}$. With $e \in E[t, t']$ we associate the following partial bijection on the vertex set of $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$, with domain and image as indicated:

$$\begin{array}{rcl} g_e \colon G_{*,t} & \longrightarrow & G_{*,t'} \\ (s, [\hat{w}]_{\hat{\mathbb{G}}}) & \longmapsto & (s, [\hat{w}\hat{e}]_{\hat{\mathbb{G}}}) = (s, [\hat{w}]_{\hat{\mathbb{G}}} \cdot \{e\} \cdot \{e, e^{-1}\} \cdot \{e^{-1}\}), \end{array}$$

where $w \in E^*[\mathbb{I}, s, t]$ and $we \in E^*[\mathbb{I}, s, t']$. Concatenation (and reduction) of corresponding words or walks in $\hat{\mathbb{I}}$ induces a well-defined groupoid operation according to

$$\begin{array}{rcl} \cdot \colon \; G_{s,t} \times G_{t,u} & \longrightarrow & G_{s,u} \\ ((s, [\hat{w}_1]_{\hat{\mathbb{G}}}), (t, [\hat{w}_2]_{\hat{\mathbb{G}}})) & \longmapsto & (s, [\hat{w}_1 \hat{w}_2]_{\hat{\mathbb{G}}}), \end{array}$$

where the concatenation relies on the condition that $\iota_2(w_1) = t = \iota_1(w_2)$. The neutral element in $G_{s,s}$ is $1_s := (s, [\lambda]_{\hat{\mathbb{G}}})$. With these stipulations,

$$\mathbb{G} := \hat{\mathbb{G}}/\mathbb{I} = (G, (G_{s,t})_{s,t \in S}, \cdot, (1_s)_{s \in S}, (g_e)_{e \in E})$$

becomes an $\mathbb{I}$-groupoid with generators

$$g_e := [e]_{\mathbb{G}} := (\iota_1(e), [\hat{e}]_{\hat{\mathbb{G}}}) \in G_{\iota_1(e), \iota_2(e)}.$$

The induced homomorphism from the free $\mathbb{I}$-groupoid (cf. Observation 8.4) onto $\mathbb{G}$ maps

$$w \in E^*[\mathbb{I}, s, t] \longmapsto [w]_{\mathbb{G}} := (\iota_1(w), [\hat{w}]_{\hat{\mathbb{G}}}) \in G_{s,t}.$$

For further analysis we also isolate the induced subgraph on those connected components of the $\hat{\mathsf{E}}$-graph $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ that embed $\mathbb{G}$:

$$\begin{array}{l} \hat{\mathbb{H}}_0 = (\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}) \restriction V_0 \subseteq \hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}} \\ \qquad \text{where } V_0 := \{(s, [u]_{\hat{\mathbb{G}}}) \in \hat{S} \times \hat{\mathbb{G}} \colon s \in S, u \in \hat{E}^*[\hat{\mathbb{I}}, s]\}. \end{array}$$

The set V_0 is the vertex set of the union of the connected components of the vertices $(s, 1)$ in $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ (i.e. of the neutral elements $(s, 1_s) \in \mathbb{G}$). Restricting further to vertices of $G \subseteq V_0$ and linking two such vertices by an e-edge if, and only if, they are linked by an $\hat{e} = \{e\}\{e, e^{-1}\}\{e^{-1}\}$-labelled walk of length 3 in $\hat{\mathbb{H}}_0 \subseteq \hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$, we obtain an $\mathbb{I}$-graph $\mathbb{H}_0$ that is interpreted in the $\hat{\mathsf{E}}$-graph $\hat{\mathbb{H}}_0$. This $\mathbb{I}$-graph $\mathbb{H}_0$ is (isomorphic to) the Cayley graph of the $\mathbb{I}$-groupoid $\mathbb{G}$:

$$\begin{array}{l} \mathbb{H}_0 = \big(G, (G_{*,s})_{s \in S}, (R_e)_{e \in E}\big) \quad \text{ where, for } e = (s, s'), \\ R_e = \big\{((s, [u]_{\hat{\mathbb{G}}}), (s', [u\{e\}\{e, e^{-1}\}\{e^{-1}\}]_{\hat{\mathbb{G}}})) \colon u \in \hat{E}^*[\hat{\mathbb{I}}, s]\big\}. \end{array}$$

Observation 9.1. *Let $\hat{\mathbb{G}}$ be an $\hat{\mathsf{E}}$-group that is compatible with $\hat{\mathbb{I}}$. Then the Cayley graph $\mathbb{C}\mathbb{G}$ of the $\mathbb{I}$-groupoid $\mathbb{G} = \hat{\mathbb{G}}/\mathbb{I}$ as just constructed from $\mathbb{C}\hat{\mathbb{G}}$ is (isomorphic to) the $\mathbb{I}$-graph $\mathbb{H}_0$ interpreted in $\hat{\mathbb{H}}_0 \subseteq \hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$.*

9.2 Transfer of acyclicity, compatibility and symmetries

The following is the main technical result of this section. It reduces the construction of N-coset acyclic groupoids to the construction of Cayley groups with involutive generators that are N-acyclic over some constraint graph.

Proposition 9.2. *For a constraint pattern $\mathbb{I}$ and its translation into a constraint graph $\hat{\mathbb{I}}$ as above, let $\hat{\mathbb{G}}$ be an $\hat{\mathsf{E}}$-groupoid that is compatible with $\hat{\mathbb{I}}$. Let $\mathbb{G} := \hat{\mathbb{G}}/\mathbb{I}$ the $\mathbb{I}$-groupoid whose Cayley graph $\mathbb{C}\mathbb{G}$ is realised as $\mathbb{H}_0$ within $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ as discussed above.*

- *(i) If $\hat{\mathbb{G}}$ is N-acyclic over the constraint graph $\hat{\mathbb{I}}$, then $\mathbb{G}$ is N-acyclic.*
- *(ii) If $\hat{\mathbb{G}}$ is compatible with the $\hat{\mathsf{E}}$-translation of a complete $\mathbb{I}$-graph $\mathbb{H}$, then $\mathbb{G}$ is compatible with $\mathbb{H}$.*
- *(iii) Any symmetry ρ of $\mathbb{I}$ induces a permutation $\hat{\rho} \in \mathrm{Sym}(\hat{E})$ that is a symmetry of $\hat{\mathbb{I}}$; if $\hat{\rho}$ is a symmetry of $\hat{\mathbb{G}}$ then ρ is a symmetry of $\mathbb{G}$.*

The main claim, concerning N-acyclicity, follows directly from the following compatibility of the corresponding notions of cycles with the interpretation of $\mathbb{C}\mathbb{G} \simeq \mathbb{H}_0$ in $\hat{\mathbb{H}}_0 \subseteq \hat{\mathbb{I}} \otimes \hat{\mathbb{G}}$. This is expressed in the following lemma; the arguments towards compatibility with a given $\mathbb{H}$ and compatibility of the whole construction with symmetries are straightforward.

Lemma 9.3. *In the situation of Proposition 9.2 there is a natural translation of coset cycles in the $\mathbb{I}$-groupoid $\mathbb{G} = \hat{\mathbb{G}}/\mathbb{I}$ based on the map $\hat{}$ for generator sets, which translates coset cycles in the groupoid $\mathbb{G}$ into $\hat{\mathbb{I}}$-coset cycles of the same length in the group $\hat{\mathbb{G}}$.*

Proof. Let

$$(*) \qquad (g_i\mathbb{G}[\alpha_i], g_i)_{i\in\mathbb{Z}_n}$$

be a coset cycle in the groupoid $\mathbb{G}$, according to Definition 8.12, viewed in $\mathbb{H}_0$. The connectivity condition for the cycle $(*)$ and the manner in which $\mathbb{H}_0$ is interpreted in $\hat{\mathbb{H}}_0 \subseteq \hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ implies that there is an $\hat{\alpha}_i$-walk from $\hat{g}_i = [\hat{w}_i]_{\hat{\mathbb{G}}}$ to $\hat{g}_{i+1} = [\hat{w}_{i+1}]_{\hat{\mathbb{G}}}$, labelled by the $\hat{}$-translation of an α_i-word of generators representing $g_i^{-1}g_{i+1} \in \mathbb{G}$. The natural $\hat{}$-translation of the cycle $(*)$ into $\hat{\mathbb{G}}$ is

$$(**) \qquad (\hat{\mathbb{G}}[\hat{\mathbb{I}}, \hat{\alpha}_i, s_i; \hat{g}_i], \hat{g}_i)_{i\in\mathbb{Z}_n},$$

where the labels $s_i \in S \subseteq \hat{S}$ are determined by the sorts of the g_i according to $s_i = \iota_2(w_i)$. This translation in effect replaces the subsets $g_i\mathbb{G}[\alpha_i]$ by their closures $\hat{\mathbb{G}}[\hat{\mathbb{I}}, \hat{\alpha}_i, s_i; \hat{g}_i]$ w.r.t. $\hat{\mathbb{I}}$-reachability inside their $\hat{\alpha}_i$-coset. This closure is obtained as the union of all $\{\{e\}, \{e, e^{-1}\}, \{e^{-1}\}\}$-connected components in

$\hat{\mathbb{H}}_0 \subseteq \hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ that contain at least one element of the representation of $g_i\mathbb{G}[\alpha_i]$ in $\mathbb{H}_0$.

It is clear that $(**)$ has the format of a potential $\hat{\mathbb{I}}$-coset cycle of length n over $\hat{\mathbb{I}}$ in $\hat{\mathbb{G}}$ in the sense of Definition 6.8. It remains to show that the separation condition in Definition 6.8 for $(**)$ follows from the analogous condition in Definition 8.12 for $(*)$.

Suppose that, in violation of the separation condition for $(**)$,

$$(\dagger) \qquad \hat{g} \in \hat{g}_i\hat{\mathbb{G}}[\hat{\mathbb{I}}, \hat{\alpha}_{i-1,i}, s_i] \cap \hat{g}_{i+1}\hat{\mathbb{G}}[\hat{\mathbb{I}}, \hat{\alpha}_{i,i+1}, s_{i+1}],$$

where $\alpha_{i,i\pm1} = \alpha_i \cap \alpha_{i\pm1}$.

We analyse this non-trivial intersection in terms of the representation of G in $\hat{\mathbb{H}}_0 \subseteq \hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$. Let $\hat{g} = \hat{g}_i[w_{i,i-1}]_{\hat{\mathbb{G}}} = \hat{g}_{i+1}[w_{i,i+1}]_{\hat{\mathbb{G}}}$ for suitable $w_{i,j} \in \hat{\alpha}^*_{i,j}[\hat{\mathbb{I}}, s_j]$. By the separation condition for $(*)$, $\hat{g}$ is not represented as an element of $\mathbb{G}$ or $\mathbb{H}_0$ in $\hat{\mathbb{H}}_0$, so that $\iota_2(w_{i,j}) \in \hat{S} \setminus S$, i.e. $\iota_2(w_{i,j}) \in \{s_e, s_{e^{-1}}\}$ for some $e \in E$.

But in $\{\{e\}, \{e, e^{-1}\}, \{e^{-1}\}\}$-components of elements of G or vertices of $\mathbb{H}_0$ in $\hat{\mathbb{H}}_0$, any vertex with ι_2-value outside S is isolated from all vertices with ι_2-value in S by $\{e\}$- and $\{e^{-1}\}$-edges (just as vertices in $\hat{S} \setminus S$ are isolated from S in $\hat{\mathbb{I}}$). So $(\dagger)$ implies that $e, e^{-1} \in \alpha_{i,j}$ for $j = i \pm 1$. This would imply that there also is an e-link between the elements of that component that represent elements of G. So elements of $g_i\mathbb{G}[\alpha_{i,i-1}]$ and of $g_{i+1}\mathbb{G}[\alpha_{i,i+1}]$ occur in the same $\{e, e^{-1}\}$-component, which would violate the separation condition for $(*)$ since $e \in \alpha_{i,i-1}$ and $e \in \alpha_{i,i+1}$. ☐

Theorem 9.4. *For every finite constraint pattern $\mathbb{I} = (S, E)$, every complete $\mathbb{I}$-graph $\mathbb{H}$ and every $N \geqslant 2$ there is a finite N-acyclic $\mathbb{I}$-groupoid $\mathbb{G}$ that is compatible with $\mathbb{H}$. Such $\mathbb{G}$ can be chosen to be fully symmetric w.r.t. the given data, i.e. such that every symmetry ρ of $\mathbb{I}$ that induces a symmetry of the $\mathbb{I}$-graph $\mathbb{H}$ is also a symmetry of the $\mathbb{I}$-groupoid $\mathbb{G}$: $\mathbb{H}^\rho \simeq \mathbb{H} \Rightarrow \mathbb{G}^\rho \simeq \mathbb{G}$.*

Choosing the Cayley graph of a given $\mathbb{I}$-groupoid $\mathbb{G}_0$ for $\mathbb{H}$, we obtain a fully symmetric N-acyclic $\mathbb{I}$-groupoid $\mathbb{G} \succcurlyeq \mathbb{G}_0$. For $\mathbb{H} := \mathrm{sym}(\mathbb{I})$ (regarding $\mathbb{I}$ as a complete $\mathbb{I}$-graph according to Definition 8.8) one obtains N-acyclic $\mathbb{I}$-groupoids that are fully symmetric over $\mathbb{I}$.

10 Conclusion and primary applications

The generic constructions of the preceding chapters show the versatility of the fruitful idea to go back and forth between group-like structures (monoids and groups as well as groupoids) and graph-like structures (graphs and multi-graphs, undirected as well as directed, and possibly vertex- or edge-coloured). In one direction the passage involves the familiar encoding of algebraic structures in the graph-like representation of generators, as in the classical notion of Cayley graphs for groups; in the converse direction, permutation groups are induced by various operations on graph-like structures. We have here tried to contribute to these connections with a special emphasis on strong algebraic-combinatorial criteria of graded acyclicity in finite structures. The constructions presented here extend techniques for the construction of N-acyclic groups with involutive generators from [13] to yield a conceptual improvement and correction of the proposed constructions for groupoids from [14]. Due to its symmetry preserving generic character, the new presentation also supports the use of these groupoids in [15] where symmetry considerations are of the essence towards lifting local symmetries to global symmetries in finite structures. In a different direction, 2-acyclic finite groupoids have also been used to resolve an open problem of a purely semigroup-theoretic nature in Bitterlich [5].

To conclude the present treatment we briefly look at the most salient application for finite groups and groupoids of graded coset-acyclicity. This concerns the construction of finite coverings of graphs and hypergraphs that unravel short cycles.

(1) Natural, unbranched finite coverings of graphs by graphs with interesting acyclicity properties can be obtained as weak subgraphs of the Cayley graphs of suitable E-groups where E is the set of edges of the graph to be covered (individually labelled as it were). While similar constructions have been used in [12, 13] and a precursor for special graphs in [8], we illustrate the key to the new generalisation in Proposition 10.1 below.

(2) Natural reduced products with N-acyclic $\mathbb{I}$-groupoids yield finite branched N-acyclic coverings of hypergraphs where a constraint pattern $\mathbb{I} = (S, E)$ is induced by the intersection graph of (V, S) that encodes the intersection pattern between hyperedges in the given hypergraph (cf. [14]).

(3) A new and more direct approach to finite branched N-acyclic coverings of hypergraphs (V, S) can be based on $\mathbb{I}$-products between a constraint graph $\mathbb{I} = (S, E)$ induced by the intersection graph of (V, S) and suitable E-groups that are not just N-acyclic but N-acyclic over $\mathbb{I}$; cf. Proposition 10.2 below.

Of these fundamental applications, (2) has been explored in stages in [13, 14, 15]. Application (1) is new in its strong form that involves the new notion of N-acyclicity of groups over a constraint graph $\mathbb{I}$. Application (3) similarly supersedes (2). Recall from Section 6 how control of cyclic configurations can be extended to configurations governed by reachability patterns w.r.t. a given constraint graph $\mathbb{I}$. While we have seen in Section 8 how such groups can yield coset acyclicity in groupoids as used in (2), the underlying groups can also be put to use directly in (1) and (3).

Graph coverings. For a finite simple graph $\mathbb{V} = (V, E)$ consider, as a set E of involutive generators for E-groups, the set of all edges $e = \{v, v'\} \in E$, and as a constraint graph $\mathbb{I}$ the E-graph $\mathbb{I} = (V, (\{e\})_{e \in E})$ ($\mathbb{V}$ with individually labelled edges). For any E-group $\mathbb{G}$ that is compatible with $\mathbb{I}$ consider the direct product $\hat{\mathbb{V}} = \mathbb{I} \otimes \mathbb{C}\mathbb{G}$ of the constraint graph $\mathbb{I}$ with the Cayley graph $\mathbb{C}\mathbb{G}$ of $\mathbb{G}$ according to Definition 6.5. Then the natural projection

$$\begin{array}{rccc} \pi\colon & \hat{\mathbb{V}} & \longrightarrow & \mathbb{V} \\ & (v,g) & \longmapsto & v \end{array}$$

provides an unbranched covering of $\mathbb{V}$ by $\hat{\mathbb{V}} = \mathbb{I} \otimes \mathbb{C}\mathbb{G}$. Recall that the connected components of $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ are isomorphic to weak subgraphs of $\mathbb{C}\mathbb{G}$ (cf. Remark 6.6).

Proposition 10.1. *Let $\mathbb{V} = (V, E)$ be a connected finite simple graph, E associated with its edge set E as above and $\mathbb{G}$ an E-group that is compatible with the E-graph $\mathbb{I} := (V, (\{e\})_{e \in E})$. Then each connected component $\mathbb{H}$ of the direct product $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$,*

(i) is realised as a weak subgraph of the Cayley graph $\mathbb{C}\mathbb{G}$ of $\mathbb{G}$ and

(ii) is an unbranched finite covering w.r.t. the natural projection $\pi\colon (v,g) \mapsto v$.

This covering graph $\mathbb{H}$ inherits the acyclicity properties of $\mathbb{C}\mathbb{G}$: if $\mathbb{G}$ is N-acyclic over $\mathbb{I}$, then $\mathbb{H}$ admits no cyclic configurations of length up to N of overlapping α_i-connected components with the natural separation condition for subsets $\alpha_i \subseteq E$.

Hypergraph coverings. With a finite hypergraph $\mathbb{V} = (V, S)$ with $S \subseteq \mathcal{P}(V)$ associate its *intersection graph* $\mathbb{I} = (S, E)$ where

$$E = \{\{s, s'\} \in S^2 \colon s \neq s', s \cap s' \neq \emptyset\}.$$

If $\mathbb{G}$ is an E-group that is compatible with $\mathbb{I}$ then the direct product $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ of the intersection graph $\mathbb{I}$ with the Cayley graph $\mathbb{C}\mathbb{G}$ of $\mathbb{G}$ gives rise to a finite branched hypergraph covering $\mathbb{V} = (V, S)$ as follows. Consider the following disjoint union of G-tagged copies of the hyperedges of $\mathbb{V}$,

$$\bigcup_{s \in S} s \times G$$

and its quotient w.r.t. the equivalence relation $\approx$ induced by identifications

$$(v, g) \approx (v, ge) \ \text{ for } \ e = \{s, s'\} \in E, v \in s \cap s'.$$

The induced equivalence is such that $(v, g) \approx (v', g')$ if, and only if, $v' = v$ and $g^{-1}g' \in \mathbb{G}[\alpha]$ for $\alpha = \alpha_v := \{e = \{s, s'\} \in E \colon v \in s \cap s'\}$.

Writing $[(v, g)]$ for the equivalence class of $(v, g) \in s \times G$, we extend this notation to the subsets induced by the $s \in S$:

$$[s, g] := \{[v, g] \colon v \in s\} \text{ for } (s, g) \in \mathbb{I} \otimes \mathbb{C}\mathbb{G}.$$

In the $\approx$-quotient, the e-edge between (s, g) and (s', ge) in $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ becomes an intersection of the copies $[s, g]$ and $[s', ge]$ of the hyperedges s and s' in the covering hypergraph. This covering hypergraph is $\hat{\mathbb{V}} := \mathbb{V} \otimes \mathbb{C}\mathbb{G} = (\hat{V}, \hat{S})$ where

$$\begin{aligned} \hat{V} &:= \{[(v, g)] \colon s \in S, g \in G\} \\ \hat{S} &:= \{[s, g] \colon s \in S, g \in G\} \end{aligned}$$

with covering projection

$$\begin{array}{rcl}\pi\colon \hat{\mathbb{V}} = (\hat{V}, \hat{S}) & \longrightarrow & \mathbb{V} = (V, S) \\ [(v,g)] & \longmapsto & v.\end{array}$$

Proposition 10.2. *Let (V, S) be a finite hypergraph, $\mathbb{I} = (S, E)$ its intersection graph. If $\mathbb{G}$ is an E-group that is compatible with $\mathbb{I}$ then the hypergraph $\hat{\mathbb{V}} := \mathbb{V} \otimes \mathbb{C}\mathbb{G}$, which is based on the $\mathbb{I}$-product $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ of $\mathbb{I}$ with the Cayley graph $\mathbb{C}\mathbb{G}$ of $\mathbb{G}$, gives rise to a finite branched hypergraph covering $\pi\colon \hat{\mathbb{V}} \longrightarrow \mathbb{V}$. This covering hypergraph $\hat{\mathbb{V}}$ inherits the acyclicity properties of $\mathbb{C}\mathbb{G}$ in the following sense: if $\mathbb{G}$ is N-acyclic over $\mathbb{I}$, then every induced sub-hypergraph on up to N vertices is acyclic in the sense of classical hypergraph theory.*

For acyclicity in hypergraph terminology (conformality and chordality and tree-decomposability), compare [2, 3].

Proof. Consider the hypergraph $\mathbb{V} \otimes \mathbb{C}\mathbb{G}$ as defined above, for an E-group $\mathbb{G}$ that is compatible with the intersection graph $\mathbb{I} = (S, E)$ of $\mathbb{V}$.

Note that in $\hat{\mathbb{V}}$, $\hat{v} \in [t, g] \cap [t', g']$ if, and only if, $\hat{v} = [(v, g)] = [(v, g')]$ for some $v \in t \cap t'$ and g, g' such that $g^{-1}g' = [w]_{\mathbb{G}}$ for some $w \in \alpha^*[\mathbb{I}, t, t']$ where $\alpha = \alpha_v = \{e = \{s, s'\} \in E \colon v \in s \cap s'\}$.

It remains to argue for N-acyclicity of $\hat{\mathbb{V}}$ if $\mathbb{G}$ is chosen to be N-acyclic over $\mathbb{I}$. We show that in this situation the Gaifman graph of $\hat{\mathbb{V}}$ cannot have chordless cycles of lengths n for $3 < n \leqslant N$ (N-chordality), nor can it have cliques of size up to N that are not contained in a single hyperedge (N-conformality).

N-chordality. Suppose $(\hat{v}_i)_{i \in \mathbb{Z}_n}$ is a chordless cycle of length $n > 3$ in the Gaifman graph of $\hat{\mathbb{V}} = (\hat{V}, \hat{S})$, and let $[s_i, g_i] \in \hat{S}$ be such that $\hat{v}_i \in [s_i, g_i] \cap [s_{i+1}, g_{i+1}]$. This implies that $\hat{v}_i$ can be represented as $\hat{v}_i = [(v_i, g_i)] = [(v_i, g_{i+1})]$ for some $v_i \in s_i \cap s_{i+1}$ and that $h_i := g_i^{-1} g_{i+1} = [w_i]_{\mathbb{G}}$ for some $w_i \in \alpha_i^*[\mathbb{I}, s_i, s_{i+1}]$ where $\alpha_i = \{e = \{s, s'\} \in E \colon v_i \in s \cap s'\}$. We claim that

$$(\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i; g_i], g_i)_{i \in \mathbb{Z}_n}$$

is an $\mathbb{I}$-coset cycle in $\mathbb{G}$, in the sense of Definition 6.8. Then $n > N$ follows from N-acyclicity of $\mathbb{G}$ over $\mathbb{I}$. Of the two conditions in Definition 6.8, connectivity is obvious; it remains to check the separation condition:

$$\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_{i,i-1}, s_i; g_i] \cap \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_{i,i+1}, s_{i+1}; g_{i+1}] = \emptyset,$$

where $\alpha_{i,j} := \alpha_i \cap \alpha_j$. This follows from chordlessness of the given cycle. Suppose g were a member of this intersection, i.e. $h := g_i^{-1} g = [w]_{\mathbb{G}}$ for some $w \in (\alpha_{i-1} \cap \alpha_i)^*[\mathbb{I}, s_i, s]$ and $h' := g_{i+1}^{-1} g = [w']_{\mathbb{G}}$ for some $w' \in (\alpha_{i+1} \cap \alpha_i)^*[\mathbb{I}, s_{i+1}, s]$ (the same s, due to compatibility of $\mathbb{G}$ with $\mathbb{I}$). Then $\hat{v}_{i-1} = [(v_{i-1}, g_{i-1})] = [(v_{i-1}, g)]$ because $w \in \alpha_{i-1}^*$ and $\hat{v}_{i-1} \in [s_i, g_i]$, which implies $\hat{v}_{i-1} \in [s, g]$. Similarly, $\hat{v}_{i+1} = [(v_{i+1}, g_{i+1})] = [(v_{i+1}, g)]$ because $w' \in \alpha_{i+1}^*$, which implies that $\hat{v}_{i+1} \in [s, g]$, too. So the given cycle would have a chord linking $\hat{v}_{i-1}$ to $\hat{v}_{i+1}$.

N-conformality. Suppose $m = \{\hat{v}_i \colon 1 \leqslant i \leqslant n\}$ forms a clique of size n in the Gaifman graph of $\hat{\mathbb{V}} = (\hat{V}, \hat{S})$ such that every subset $m_i := m \setminus \{\hat{v}_{i-1}\}$ of size $n-1$ is contained in some hyperedge (a minimal violation of conformality). Let $\hat{v}_i = [(v_i, g_i)]$, $h_i := g_i^{-1} g_{i+1}$. For $1 \leqslant i \leqslant n$, let $[s_i, g_i] \in \hat{S}$ be a hyperedge that contains $m_i = m \setminus \{[(v_{i-1}, g_{i-1})]\}$. Therefore $\hat{v}_j = [(v_j, g_j)] \in [s_i, g_i]$ for all $j \neq i-1$. Let $\alpha_i = \{e = \{s, s'\} \in E \colon v_i \in s \cap s'\}$ and put $\beta_i := \bigcap_{j \neq i-1} \alpha_j$ so that $\hat{v} = [(v, g_j)] = [(v, g)]$ for all $g \in \mathbb{CG}[\mathbb{I}, \beta_i, s_j; g_j]$, $v \in m_i$ and $j \neq i-1$. Note that any intersection $\beta_{i,j} := \beta_i \cap \beta_j$ for $i \neq j$ is just $\bigcap_{i \in \mathbb{Z}_n} \alpha_i =: \beta$. Consider

$$(\mathbb{CG}[\mathbb{I}, \beta_i, s_i; g_i], g_i)_{i \in \mathbb{Z}_n}$$

as a candidate for an $\mathbb{I}$-coset cycle. We show that if this is not an $\mathbb{I}$-coset cycle, then the whole of m is contained in some hyperedge $[s, g] \in \hat{S}$. Again, the connectivity condition on $\mathbb{I}$-coset cycles from Definition 6.8 is obvious for the given data. The separation condition now is that

$$\mathbb{CG}[\mathbb{I}, \beta_{i,i-1}, s_i; g_i] \cap \mathbb{CG}[\mathbb{I}, s_{i+1}, \beta_{i,i+1}, s_{i+1}; g_{i+1}] = \emptyset,$$

where $\beta_{i,i-1} = \beta_{i,i+1} = \beta = \bigcap_{i \in \mathbb{Z}_n} \alpha_i$. Assume there were some g in this intersection, i.e. $g = g_i h$ for some $h = [w]_{\mathbb{G}}$ with $w \in \beta^*[\mathbb{I}, s_i, s]$ and $g = g_{i+1} h'$ for some $h' = [w']_{\mathbb{G}}$ with $w' \in \beta^*[\mathbb{I}, s_{i+1}, s]$. We claim that this would imply $m \subseteq [s, g]$. This follows as $\hat{v}_j = [(v_j, g_j)] = [(v_j, g)] \in [s, g]$ for $j \neq i-1$, by the nature of $h = [w]_{\mathbb{G}}$ and since $\hat{v}_j \in [s_i, g_i]$, and as $\hat{v}_j = [(v_j, g_j)] = [(v_j, g)] \in [s, g]$ for $j \neq i$, by the nature of $h' = [w']_{\mathbb{G}}$ and since $\hat{v}_j \in [s_{i+1}, g_{i+1}]$. ☐

Acknowledgements. This paper has a complicated history. It is based on ideas revolving about acyclicity in hypergraphs and Cayley groups that I first expounded in [13] (the journal version of a LICS 2010 paper). That paper deals with the construction of finite N-acyclic groups and applies it to the finite model theory of the guarded fragment. The promising extension to the groupoid situation was seemingly achieved with [14], which also paved the way to more fundamental applications in hypergraph coverings and extension properties for partial automorphisms. These applications were successively elaborated in a series of arXive preprints leading to [15] and stand as key motivations and achievements also of my DFG project on *Constructions and Analysis in Hypergraphs of Controlled Acyclicity*, 2013–18. The proposed construction of N-acyclic groupoids in [14], however, contained a serious flaw that was also carried in the arXive preprints. This mistake from [14] was finally found out in 2019 by Julian Bitterlich *after* he had shown, as part of his PhD work, that those results would positively resolve a longstanding conjecture by Henkell and Rhodes in semigroup theory, which had attracted the attention of, among others, Karl Auinger. I am deeply indebted to Julian Bitterlich for identifying and clarifying the gap in my earlier attempts at the construction of finite N-acyclic groupoids. Initial concerns about the use of my sketchy and, as it later turned out, truly flawed construction from [14] in Julian Bitterlich's contributions had earlier been raised by Jiri Kadourek. In a way, it was his initial challenge of the result that triggered the intense re-investigation that led to Julian's discovery of the actual flaw. I am also very grateful to Karl Auinger and Julian Bitterlich for helpful discussions in these difficult matters at the time of Julian's breakthrough and discovery of my mistake.

References

[1] N. Alon. Tools from higher algebra. In R. Graham, M. Grötschel, and Lovász, editors, *Handbook of Combinatorics*, volume II, pp. 1749–1783. North-Holland, 1995.

[2] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30, pp. 497–513, 1983.

[3] C. Berge. *Graphs and Hypergraphs*. North-Holland, 1973.

[4] N. Biggs. Cubic graphs with large girth. In G. Blum et al., editor, *Annals of the New York Academy of Sciences*, vol. 555, pp. 56–62. 1989.

[5] J. Bitterlich. Finite F-inverse covers do exist. Preprint *arXiv*:1808.05161, 2018.

[6] A. Cayley. Desiderata and Suggestions: No. 1. The Theory of Groups. *American Journal of Mathematics*, vol. 1, pp. 50–52. 1878.

[7] A. Cayley. Desiderata and Suggestions: No. 2. The Theory of Groups. *American Journal of Mathematics*, vol. 1, pp. 174–176. 1878.

[8] F. Canavoi and M. Otto. Common knowledge & multi-scale locality analysis in Cayley structures. In *Proceedings of 32nd Annual ACM/IEEE Symposium on Logic in Computer Science LICS'17*, 2017.

[9] R. Fox. Covering spaces with singularities. In *A Symposium in Honor of S. Lefschetz*, pp. 243–257. Princeton University Press, 1957.

[10] R. Frucht. Herstellung von Graphen mit vorgegebener abstrakter Gruppe. *Compositio Mathematica*, 6, pp. 239–250, 1938.

[11] M. Lawson. *Inverse Semigroups*. World Scientific, 1998.

[12] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130, pp. 173–205, 2004.

[13] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59 (1), 2012.

[14] M. Otto. Groupoids, hypergraphs and symmetries in finite models. In *Proceedings of 28th Annual IEEE Symposium on Logic in Computer Science LICS'13*, 2013.

[15] M. Otto. Amalgamation and symmetry: from local to global consistency in the finite. Technical report, 2018. See *arXiv*: 1709.00031.