

Acyclicity in Finite Groups and Groupoids

Martin Otto*
Department of Mathematics
Technische Universität Darmstadt

October 2022

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.

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

Contents

1	Introduction	3
2	General patterns	6
2.1	Cayley & Biggs: the basic construction	6
2.2	E-graphs and E-groups	8
2.3	Compatibility and homomorphisms	11
3	Coset cycles and acyclicity criteria	15
4	Free amalgams of E-graphs	17
4.1	Amalgamation chains	20
4.2	Amalgamation clusters	21
5	Construction of N-acyclic E-groups	24
6	Constraints on generator sequences	26
7	Small coset amalgams over I-skeletons	31
8	Construction of N-acyclic E-groups over I	42
9	From groups to groupoids	44
9.1	Constraint patterns for groupoids	44
9.2	I-groupoids and their Cayley graphs	46
9.3	Coset acyclicity for groupoids	49
10	Construction of N-acyclic I-groupoids	49
10.1	Groupoids from groups	51
10.2	Transfer of acyclicity, compatibility and symmetries	52
11	Conclusion and primary applications	54

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 [11] 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 [7, 8]. 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 from 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 [5] 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 [13] and [14] 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 [10]). 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 [15, 16].

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 [15, 16]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 [5, 13], 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 11.1 and 11.3.

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 9 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

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

edge-coloured graphs. In some instances, and especially in Sections 9 and 10, 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

We implicitly always assume that structures need only be identified up to isomorphism. This will often allow us to avoid notational complications in the interest of clarity. For instance, it often makes sense to suppress explicit notation for isomorphic embeddings if instead we can w.l.o.g. treat the pre-image as an actual (weak) substructure of the target structure; relational structures of particular interest, especially Cayley graphs, will be homogeneous in the sense that any two elements are linked by an automorphism, which implies that, up to isomorphism, explicit choice of a distinguished element can be suppressed and only the relative position of two or more elements gives rise to meaningful distinctions.

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, \dots, 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

²Depending 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).

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 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 [7, 8]. 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 [11]. 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 [5] and outlined in [1].

Let E be a finite set of letters, $|E| = d \geq 2$, to be used to label involutive generators of a group to be constructed. With E and a parameter $n \geq 1$ 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* $w \in E^{\leq n} \subseteq E^*$, i.e. strings $w = e_1 \cdots e_m$ of length $|w| = m$, $0 \leq m \leq n$, with $e_i \in E$ for $1 \leq i \leq m$ and $e_{i+1} \neq e_i$ for $1 \leq i < m$. 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

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

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

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

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 $\pi_e \in \text{Sym}(V)$ 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 $\text{Sym}(V)$ generated by these involutions:

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

For the group operation we use the convention that the action by the generators is regarded as a right action via composition, i.e. for $\rho \in G$:

$$\begin{aligned} \rho\pi_e = \pi_e \circ \rho : V &\longrightarrow V \\ w &\longmapsto \pi_e(\rho(w)). \end{aligned}$$

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

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

one for each $e \in E$. Here these edge relations are symmetric due to the involutive nature of the π_e in $\text{Sym}(V)$, and they are irreflexive and pairwise disjoint since $\text{id}_V \neq \pi_e \neq \pi_{e'}$ 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:

$$\begin{aligned} h: G &\longrightarrow G \\ g &\longmapsto hg, \end{aligned}$$

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 $k \geq 1$ 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 $\text{Sym}(V)$.

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 $\text{dom}(R)$ and range $\text{rng}(R)$ 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 $\bar{R}_e := R_e \cup \{(v, v) : v \in V \setminus \text{dom}(R_e)\}$.

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{aligned} [\]_{\mathbb{G}} : 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{aligned}$$

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 $\text{sym}(\mathbb{H})$ be the subgroup of $\text{Sym}(V)$ that is generated by the involutive permutations $\pi_e : V \rightarrow V$ 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 $\text{sym}(\mathbb{H})$ 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 $\text{sym}(\mathbb{H})$. 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 $\text{sym}(\mathbb{T}(E, n))$.

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

Recall that we let permutations act from the right. In terms of the group product in $\text{sym}(\mathbb{H})$ 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{aligned} [\]_{\mathbb{H}}: E^* &\longrightarrow \text{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{aligned}$$

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

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 $\mathbb{C}\mathbb{G} := \text{Cayley}(\mathbb{G}, E) = (G, (R_e)_{e \in E})$ with vertex set G and edge sets

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

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

The Cayley graph $\mathbb{C}\mathbb{G}$ 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 $e \in G$, $\pi_e: g \mapsto ge$. 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 $\text{Sym}(G)$ 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 $\mathbb{C}\mathbb{G} = \text{Cayley}(\mathbb{G}, E)$ is a complete and strict E-graph in the sense of Definition 2.1, and

$$\mathbb{G} = (G, \cdot, 1) \simeq \text{sym}(\mathbb{C}\mathbb{G}).$$

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 $\pi_e: g \mapsto ge$ in \mathbb{G} or in its Cayley graph $\mathbb{C}\mathbb{G}$. We similarly identify the family of generators $(\pi_e)_{e \in E}$ of $\text{sym}(\mathbb{H})$ with a subset $E \subseteq \text{sym}(\mathbb{H})$ whenever $\text{sym}(\mathbb{H})$ 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

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

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

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

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 α -walk of length n from v to v' is a sequence $v_0, e_1, v_1, e_2, \dots, e_n, v_n$ 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 typically 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 $\mathbb{H} \upharpoonright \alpha(v)$.

Note that the Cayley graph $\mathbb{C}\mathbb{G}[\alpha] \subseteq_w \mathbb{C}\mathbb{G}$ of $\mathbb{G}[\alpha]$ also arises as the α -component of $1 \in G$ in the Cayley graph $\mathbb{C}\mathbb{G}$. It is also useful to note that, if $v = v_0, e_1, v_1, e_2, \dots, e_n, v_n = v'$ 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 $\text{sym}(\mathbb{H})$ on \mathbb{H} .

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 $h: \hat{\mathbb{G}} \rightarrow \mathbb{G}$ between E-groups $\hat{\mathbb{G}}$ and \mathbb{G} then it must be

$$\begin{aligned} h: \hat{\mathbb{G}} &\longrightarrow \mathbb{G} \\ [w]_{\hat{\mathbb{G}}} &\longmapsto [w]_{\mathbb{G}} \end{aligned}$$

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 $\hat{\mathbb{G}}$. We may think of $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$ as an *unfolding* of \mathbb{G} , which also captures the relationship between the associated Cayley graphs $\mathbb{C}\hat{\mathbb{G}}$ and $\mathbb{C}\mathbb{G}$.

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 $\text{sym}(\mathbb{H})$, i.e. if $\text{sym}(\mathbb{H}) \preccurlyeq \mathbb{G}$. More generally, for $\alpha \subseteq E$, $\mathbb{G}[\alpha]$ is *compatible with* the E-graph \mathbb{H} if $\mathbb{G}[\alpha]$ is compatible with the α -reduct $\mathbb{H} \upharpoonright \alpha$ of \mathbb{H} .

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

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

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

Note that trivially $\text{sym}(\mathbb{H})$ 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 $\bigoplus_i \mathbb{H}_i$ of \mathbb{E} -graphs \mathbb{H}_i if, and only if, it is compatible with each component \mathbb{H}_i .

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 \mathbb{E} -groups with specific acyclicity properties.

Lemma 2.10. Let $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$ be \mathbb{E} -groups, $\hat{\mathbb{G}} = \text{sym}(\mathbb{H})$ for an \mathbb{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 -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 $\hat{\mathbb{G}}[\alpha] = \text{sym}(\mathbb{H} \upharpoonright \alpha)$, where $\mathbb{H} \upharpoonright \alpha$ stands for the α -reduct of \mathbb{H} , which is an α -graph. If $\mathbb{G}[\alpha]$ is compatible with every connected component of $\mathbb{H} \upharpoonright \alpha$, then $[w]_{\mathbb{G}} = [u]_{\mathbb{G}}$ for $w, u \in \alpha^*$ implies that $\pi_w = \pi_u$ in $\text{sym}(\mathbb{H} \upharpoonright \alpha)$ and therefore also in $\text{sym}(\mathbb{H})$, i.e. in $\hat{\mathbb{G}}$. \square

Our constructions of finite \mathbb{E} -groups with special combinatorial properties proceed by induction on the number of generators for subgroups under consideration, i.e. by induction on the size $|\alpha|$ of generator subsets $\alpha \subseteq E$. Subsets α of the same size are always treated in parallel in order to guarantee that our constructions are fully isomorphism-preserving and do not break any symmetries (unlike a construction governed by, for instance, a chosen enumeration of the generator set E). The desired properties will successively be achieved for subgroups generated by increasing numbers of generators from E , mostly by suitable unfolding steps from \mathbb{G} to $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$, where the desired property for $\hat{\mathbb{G}}[\alpha]$ relies on the behaviour of the subgroups $\mathbb{G}[\alpha']$ of \mathbb{G} for $\alpha' \subsetneq \alpha$; or such that all generated subgroups of $\hat{\mathbb{G}}$ generated by up to k generators inherit the property at hand from the subgroups of \mathbb{G} that are generated by fewer than k generators. So, for $\alpha \subseteq E$, we let

$$\Gamma_\alpha := \{\alpha \subseteq E : \alpha' \subsetneq \alpha\},$$

and, for $1 \leq k \leq |E| + 1$,

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

We denote corresponding families of generated subgroups and their Cayley graphs in a given \mathbb{E} -group \mathbb{G} as

$$\begin{aligned} \Gamma_k(\mathbb{G}) &:= (\mathbb{G}[\alpha'] : \alpha' \in \Gamma_k) \\ \Gamma_k(\mathbb{C}\mathbb{G}) &:= (\mathbb{C}\mathbb{G}[\alpha'] : \alpha' \in \Gamma_k) \end{aligned}$$

Recall that $\mathbb{G} \preccurlyeq \hat{\mathbb{G}}$ means that there is a homomorphism from $\hat{\mathbb{G}}$ to \mathbb{G} , or that $\hat{\mathbb{G}}$ is an unfolding of \mathbb{G} .

Definition 2.11. For $\alpha \subseteq E$ and E-groups \mathbb{G} and $\hat{\mathbb{G}}$, $\mathbb{G} \preceq_\alpha \hat{\mathbb{G}}$ denotes the relationship that $\mathbb{G} \preceq \hat{\mathbb{G}}$ and

$$\Gamma_\alpha(\mathbb{G}) = \Gamma_\alpha(\hat{\mathbb{G}}),$$

i.e. that subgroups generated by proper subsets of α are preserved in the unfolding. For $1 \leq k \leq |E| + 1$, $\mathbb{G} \preceq_k \hat{\mathbb{G}}$ analogously stands for an unfolding relationship $\mathbb{G} \preceq \hat{\mathbb{G}}$ in which

$$\Gamma_k(\mathbb{G}) = \Gamma_k(\hat{\mathbb{G}}),$$

where \mathbb{G} and $\hat{\mathbb{G}}$ agree on all their subgroups generated by fewer than k generators.

If we think of the passage from \mathbb{G} to $\hat{\mathbb{G}}$ in $\mathbb{G} \preceq \hat{\mathbb{G}}$ as an *unfolding*, then $\mathbb{G} \preceq_k \hat{\mathbb{G}}$ says that the unfolding is trivial up to the level of Γ_k , or is *conservative* w.r.t. Γ_k -generated subgroups.

Recall that $\hat{\mathbb{G}} = \text{sym}(\mathbb{H} \oplus \mathbb{C}\mathbb{G})$ guarantees $\hat{\mathbb{G}} \succ \mathbb{G}$. In this situation, moreover $\hat{\mathbb{G}} \succ_k \mathbb{G}$ if already \mathbb{G} itself is compatible with all α -components of $\mathbb{H} \oplus \mathbb{C}\mathbb{G}$ for $\alpha \in \Gamma_k$. It is clear that $\mathbb{G}[\alpha]$ is compatible with $\mathbb{C}\mathbb{G}[\alpha]$ for all α ; as we shall in Section 3, compatibility of \mathbb{G} with all $\mathbb{C}\mathbb{G}[\alpha] \in \Gamma_k(\mathbb{C}\mathbb{G})$ implies that $\mathbb{C}\mathbb{G}$ does not admit 2- or 3-cycles formed by cosets generated by fewer than k generators, so that compatibility of \mathbb{G} with all $\mathbb{C}\mathbb{G}[\alpha]$ implies 3-acyclicity of \mathbb{G} , as stated in Lemma 3.5.

The following summarises the content of Observation 2.9 and Lemma 2.10 in a straightforward manner.

Lemma 2.12. *Consider a functor F that maps E-groups \mathbb{G} to E-graphs $F(\mathbb{G})$ in an isomorphism respecting manner (i.e. $F(\mathbb{G}) \simeq F(\mathbb{G}')$ whenever $\mathbb{G} \simeq \mathbb{G}'$). Then for $\hat{\mathbb{G}} := \text{sym}(F(\mathbb{G}))$:*

- (a) *If $F(\mathbb{G})$ has a component isomorphic to $\mathbb{C}\mathbb{G}$, then $\mathbb{G} \preceq \hat{\mathbb{G}}$.*
- (b) *If F is as in (a) and such that all subgroups in $\Gamma_\alpha(\mathbb{G})$ are compatible with $F(\mathbb{G})$, then $\hat{\mathbb{G}}[\alpha'] \simeq \mathbb{G}[\alpha']$ for all $\alpha' \subsetneq \alpha$, i.e. $\mathbb{G} \preceq_\alpha \hat{\mathbb{G}}$.*
- (c) *If F is as in (a) and such that all generated subgroups $\mathbb{G}[\alpha] \in \Gamma_k(\mathbb{G})$ are compatible with $F(\mathbb{G})$, then $\hat{\mathbb{G}}[\alpha] \simeq \mathbb{G}[\alpha]$ for $\alpha \in \Gamma_k$, i.e. $\mathbb{G} \preceq_k \hat{\mathbb{G}}$.*

The pre-conditions in (b) and (c) should be seen as *downward compatibility* conditions. They guarantee that the transition from \mathbb{G} to $\hat{\mathbb{G}}$ is conservative w.r.t. to smaller generated subgroups, while possibly unfolding \mathbb{G} in non-trivial ways at the level of larger generated subgroups.

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 $|\alpha|$ of the generator set of subgroups $\mathbb{G}[\alpha]$ that may form certain obstructive patterns and progresses to exclude them by replacing \mathbb{G} essentially by suitable $\hat{\mathbb{G}} := \text{sym}(\mathbb{C}\mathbb{G} \oplus \mathbb{H})$. The crux of the matter is to find \mathbb{H} that eliminates the obstructive patterns at level Γ_{k+1} while retaining the subgroups at level Γ_k , which have already been righted in previous steps.

Slightly generalising the situation of the last lemma, consider some functor F that maps families of generated subgroups $\mathbb{G}[\alpha]$ of \mathbf{E} -groups \mathbb{G} , for $\alpha \subseteq E$, to \mathbf{E} -graphs in an isomorphism respecting manner. Let $F(\Gamma_k(\mathbb{G}))$ stand for the F -images of families of Γ_k -generated subgroups of \mathbb{G} , $F(\mathbb{G})$ for their union $F(\mathbb{G}) = \bigcup_k F(\Gamma_k(\mathbb{G}))$ and assume that $F(\mathbb{G})$ is finite up to isomorphism, for every \mathbb{G} .

Definition 2.13. We say that F is *conservative* if, for $\beta \in \Gamma_{k+1}$, all β -components of \mathbf{E} -graphs in $F(\Gamma_{k+1}(\mathbb{G}))$ are (isomorphic to) $\mathbb{C}\mathbb{G}[\beta]$ or in $F(\Gamma_k(\mathbb{G}))$.

Examples of this include amalgamation chains over α_i -graphs $\mathbb{C}\mathbb{G}[\alpha_i] \subseteq_w \mathbb{C}\mathbb{G}$ (of some bounded length) to be treated in Section 4.1, and free amalgamation clusters of α_i -graphs $\mathbb{C}\mathbb{G}[\alpha_i] \subseteq_w \mathbb{C}\mathbb{G}$ to be treated in Section 4.2.

Proposition 2.14. *Let F be a conservative functor in the above sense, \mathbb{G} an \mathbf{E} -group, $1 \leq k \leq |E|+1$. If \mathbb{G} is compatible with the \mathbf{E} -graphs in $F(\Gamma_k(\mathbb{G}))$, then \mathbb{G} admits an unfolding $\hat{\mathbb{G}} \succ_k \mathbb{G}$ that is compatible with all \mathbf{E} -graphs in $F(\mathbb{G})$.*

Proof. Inductively, one obtains a sequence of unfoldings

$$\mathbb{G} = \mathbb{G}_k \preceq_k \mathbb{G}_{k+1} \preceq_{k+1} \cdots \preceq_{n-1} \mathbb{G}_n \preceq_n \mathbb{G}_{n+1} =: \hat{\mathbb{G}}$$

for $n = |E|$. Each step $\mathbb{G}_m \preceq_m \mathbb{G}_{m+1}$ is such that \mathbb{G}_{m+1} is made compatible with $F(\Gamma_{m+1}(\mathbb{G}_m))$: \mathbb{G}_k is compatible with $F(\Gamma_k(\mathbb{G}))$ by assumption, and \mathbb{G}_{m+1} is obtained as $\text{sym}(\mathbb{H}_m)$ where \mathbb{H}_m is the disjoint union over the \mathbf{E} -graphs $\mathbb{C}\mathbb{G}_m$ and $F(\Gamma_{m+1}(\mathbb{G}_m))$. On the one hand, the component $\mathbb{C}\mathbb{G}_m$ in \mathbb{H}_m guarantees $\mathbb{G}_{m+1} \succ \mathbb{G}_m$. On the other hand, the conservative nature of F makes sure that $\mathbb{G}_{m+1}[\alpha] \simeq \mathbb{G}_m[\alpha]$ is preserved for $\alpha \in \Gamma_m$, since all non-trivial α -components of \mathbf{E} -graphs in $F(\Gamma_{m+1}(\mathbb{G}_m))$ are in $F(\Gamma_m(\mathbb{G}_m))$ so that already $\mathbb{G}_m[\alpha] \subseteq \mathbb{G}_m$ is compatible with those. \square

Symmetries. Some application contexts call for an analysis of symmetries of \mathbf{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 \mathbf{E} -groups since those, as special homomorphisms, need to fix the generators individually (model-theoretically they are treated as constants). Similarly for \mathbf{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 $\rho \in \text{Sym}(E)$ induce what in model-theoretic terminology is a *renaming*, sending \mathbb{G} to \mathbb{G}^ρ and \mathbb{H} to \mathbb{H}^ρ . For instance the ρ -renaming of the \mathbf{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 2.15. [symmetry over E]

A permutation $\rho \in \text{Sym}(E)$ of the set E is a *symmetry* of an \mathbf{E} -group \mathbb{G} if the renaming of generators according to ρ yields an isomorphic \mathbf{E} -group, $\mathbb{G}^\rho \simeq \mathbb{G}$. Similarly, ρ is a symmetry of the \mathbf{E} -graph \mathbb{H} if the renaming of its edge relations according to ρ yields an isomorphic \mathbf{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 $\rho \in \text{Sym}(E)$ is a symmetry; the same is then true of the resulting E-group $\mathbb{G} = \text{sym}(\mathbb{T}(E, n))$ and its Cayley graph $\mathbb{C}\mathbb{G}$.

3 Coset cycles and acyclicity criteria

A basic notion of n -acyclicity for E-groups would just forbid 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, as the length of the shortest non-trivial generator cycle in \mathbb{G} is the length of the shortest graph cycle in $\mathbb{C}\mathbb{G}$, i.e. its girth. We are here 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 $g_i \mathbb{C}\mathbb{G}[\alpha_i]$.

Definition 3.1. [coset cycle]

Let \mathbb{G} be an E-group, $n \geq 2$. 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 $1 \leq k \leq |E|$. 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 $\mathbb{C}\mathbb{G}$ 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 $n \leq N$ 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 $|\alpha| \leq k$. 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.

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. This view also points to a natural generalisation of the concept from Cayley graphs to more general E-graphs, in Definition 4.3.

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 – another concept that will be generalised from Cayley graphs of E-groups to other classes of E-graphs in Sections 4 and 7.

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

- (i) *If \mathbb{G} is 2-acyclic then there is a unique \subseteq -minimal generator set $\alpha_0 \subseteq E$ such that $g \in \mathbb{G}[\alpha_0]$, viz. $\alpha_0 := \bigcap \{\alpha \subseteq E : g \in \mathbb{G}[\alpha]\}$.*
- (ii) *If \mathbb{G} is 3-acyclic then there is a unique \subseteq -minimal generator set $\alpha_0 \subseteq E$ such that $\mathbb{G}[\alpha_0] \cap B \neq \emptyset$, viz. $\alpha_0 := \bigcap \{\alpha \subseteq E : \mathbb{G}[\alpha_0] \cap B \neq \emptyset\}$.*

In connection with (ii) it is not hard to see that, conversely, the stated uniqueness property implies 3-acyclicity.

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 minimal supporting sets of generators α_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 $1\mathbb{G}[\alpha_1] = \mathbb{G}[\alpha_1]$, $g_1\mathbb{G}[\beta]$ 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 $g' \in g_1\mathbb{G}[\alpha_1 \cap \beta] \cap 1\mathbb{G}[\alpha_1 \cap \alpha_2]$; now $g' \in g_1\mathbb{G}[\alpha_1 \cap \beta]$ 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 $g_1\mathbb{G}[\beta]$ finally means that there is some $g' \in g_1\mathbb{G}[\alpha_1 \cap \beta] \cap g_2\mathbb{G}[\alpha_2 \cap \beta]$ which implies that $g' \in B$ as before. Moreover $g' \in \mathbb{G}[\alpha_i]$ as $g_i \in \mathbb{G}[\alpha_i]$ and $g' \in g_i\mathbb{G}[\alpha_i]$ for $i = 1, 2$; so $g' \in \mathbb{G}[\alpha_1] \cap \mathbb{G}[\alpha_2] = \mathbb{G}[\alpha_1 \cap \alpha_2] = \mathbb{G}[\alpha_0]$. \square

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 3-cycle in \mathbb{G} . The separation condition implies that $\alpha_i \subsetneq E$ and that, for instance,

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

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

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

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 deletion 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 $[w_2w_0w_1]_{\mathbb{H}} = 1 \in \text{sym}(\mathbb{H})$ 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 $\text{sym}(\mathbb{C}\mathbb{G}[\alpha_0])$ 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. \square

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 $h: \hat{\mathbb{G}}[\alpha] \rightarrow \mathbb{G}[\alpha]$ 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 $\alpha_{i,i\pm 1}$ -cosets as its subsets. \square

The method suggested by Proposition 2.14 is to be used to eliminate N -cycles of cosets with increasing numbers of generators, in an inductive treatment. The relevant configurations in $F(\mathbb{G})$ will be suitable amalgams of E-graphs $\mathbb{C}\mathbb{G}[\alpha]$ for $\alpha \in \Gamma_k$. In light of the proposition we need to focus on functors F that are conservative in the sence of Definition 2.13. Towards coset acyclicity we can use as F -images amalgamation chains that unfold potential cycles.

4 Free amalgams of E-graphs

We investigate free amalgams of E-graphs, and especially of Cayley graphs of generated subgroups $\mathbb{G}[\alpha]$. The goal is to provide elements of a structure theory for Cayley graphs that provides tools for partial unfoldings that eliminate obstructions to desirable combinatorial properties, in an inductive process that treats generated subgroups in $\Gamma_k(\mathbb{G})$ for increasing k .

The core idea of free amalgams of E-graphs is to superpose copies of suitably matching E-graphs in distinguished elements or regions with just minimal identifications. Those minimal identifications, or regions of overlap, are determined by shared generator edges. We first look at superpositions of two E-graphs, then expand the idea to certain chains or clusters of several E-graphs and in particular of E-graphs of the form $\mathbb{C}\mathbb{G}[\alpha_i]$ that arise as substructures of a common E-group \mathbb{G} . In each case, the isomorphism type of the resulting E-graphs is fully determined by the isomorphism types of the constituents – and not by the manner in which e.g. constituents $\mathbb{C}\mathbb{G}[\alpha_i]$ 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 $\mathbb{C}\mathbb{G}$ on a union of α_i -cosets of \mathbb{G} (cf. Observations 4.2, 4.6, 4.11). For instance, if \mathbb{G} is 2-acyclic, then the weak subgraphs of $\mathbb{C}\mathbb{G}$ on overlapping cosets




Figure 1: Free amalgamation.

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 $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2]$ (Observation 4.2).

Definition 4.1. [free amalgam]

A *free amalgam* between two pointed E-graphs \mathbb{H}_1, v_1 and \mathbb{H}_2, v_2 is defined relative to a designated pointed E-graph \mathbb{H}_0, v_0 , which serves as the *overlap*, if \mathbb{H}_0 is connected and admits isomorphic embeddings $\rho_i: \mathbb{H}_0, v_0 \rightarrow \mathbb{H}_i, v_i$ onto weak substructures $\rho_i(\mathbb{H}_0) \subseteq_w \mathbb{H}_i$ that map v_0 to v_i , for $i = 1, 2$, and if the following condition is satisfied for all $v \in \mathbb{H}_0$ and $e \in E$:

$$(*) \quad \left(\rho_1(v) \in \text{dom}(R_e^{\mathbb{H}_1}) \text{ and } \rho_2(v) \in \text{dom}(R_e^{\mathbb{H}_2}) \right) \Rightarrow v \in \text{dom}(R_e^{\mathbb{H}_0}).$$

In this situation we say that \mathbb{H}_1, v_1 and \mathbb{H}_2, v_2 *admit a free amalgam over* \mathbb{H}_0, v_0 , and we obtain this free amalgam as

$$\mathbb{H}_1, v_1 \oplus_{\mathbb{H}_0, v_0} \mathbb{H}_2, v_2$$

from the disjoint union of \mathbb{H}_1 and \mathbb{H}_2 by identification of vertices $\rho_1(v)$ and $\rho_2(v)$ for all $v \in \mathbb{H}_0$. When referring to this free amalgam as a pointed E-graph we regard the copy of v_0 , i.e. the identification of v_1 and v_2 , as the distinguished vertex.

We also think of the constituents \mathbb{H}_i, v_i for $i = 1, 2$ as well as of the overlap \mathbb{H}_0, v_0 as isomorphically embedded as weak substructures of the free amalgam, via the maps $\sigma_i: \mathbb{H}_i \rightarrow \mathbb{H}$ and $\sigma_i \circ \rho_i: \mathbb{H}_0 \rightarrow \mathbb{H}$, for both $i = 1$ and $i = 2$, as in the commuting diagram of Figure 1.

Condition (*) serves to guarantee that the resulting structure is an E-graph, i.e. that identifications of vertices from \mathbb{H}_1 and \mathbb{H}_2 cannot introduce forking e -edges. In all situations to be considered below, however, we shall be dealing with more restricted settings where the embeddings ρ_i are onto induced substructures $\rho_i(\mathbb{H}_0) \subseteq \mathbb{H}_i$. In most cases \mathbb{H}_0, v_0 will be (isomorphic to) the α_0 -component of v_i in \mathbb{H}_i for at least one of $i = 1, 2$. In many cases, even $\mathbb{H}_0, v_0 \simeq \mathbb{H}_1[v_1; \alpha_0] \simeq \mathbb{H}_2[v_2; \alpha_0]$.

The most basic case occurs if the \mathbb{H}_i arise as Cayley graphs of subgroups of a common E-group \mathbb{G} : $\mathbb{H}_i \simeq \mathbb{C}\mathbb{G}[\alpha_i] \subseteq_w \mathbb{C}\mathbb{G}$. In this case the isomorphism

type of the free amalgam of \mathbb{H}_1 and \mathbb{H}_2 does not depend on the choice of the distinguished vertices, and \mathbb{H}_0 is determined relative to \mathbb{H}_1 and \mathbb{H}_2 . Together with the requirement of connectivity for \mathbb{H}_0 , condition (*) implies that the overlap is $\mathbb{H}_0 \simeq \mathbb{G}[\alpha_0]$ for $\alpha_0 := \alpha_1 \cap \alpha_2$. We correspondingly denote such free amalgams as just $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2]$ whenever it is clear that we refer to an amalgam (rather than to a disjoint union). If \mathbb{G} is 2-acyclic, then

$$\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2] \simeq \mathbb{C}\mathbb{G} \upharpoonright (\alpha_1(1) \cup \alpha_2(1)) \subseteq_w \mathbb{C}\mathbb{G},$$

and this (essentially unique) embedding even yields an induced substructure relationship $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2] \subseteq \mathbb{C}\mathbb{G}$ if \mathbb{G} is 3-acyclic. For these claims observe that no two elements of $G[\alpha_i]$ can be linked by any generator $e \notin \alpha_i$ due to 2-acyclicity; similarly, 3-acyclicity rules out e -links for $e \notin \alpha_1 \cup \alpha_2$ between $G[\alpha_1] \setminus G[\alpha_2]$ and $G[\alpha_2] \setminus G[\alpha_1]$.

Observation 4.2. *For 2-acyclic \mathbb{G} and $\alpha_1, \alpha_2 \subseteq E$, the free amalgam $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2]$ embeds into $\mathbb{C}\mathbb{G}$ as a weak substructure. For 3-acyclic \mathbb{G} it embeds as an induced substructure. Moreover, \mathbb{G} is 2-acyclic if, and only if, all free amalgams of the form $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2]$ embed into $\mathbb{C}\mathbb{G}$ as weak substructures.*

Generalising 2-acyclicity from Cayley graphs of E-groups to arbitrary E-graphs, we obtain the following notion of simple connectivity in terms of α -components.

Definition 4.3. [2-acyclicity]

An E-graph \mathbb{H} is 2-acyclic if, for all $\alpha_1, \alpha_2 \subseteq E$ with intersection $\alpha_0 := \alpha_1 \cap \alpha_2$ and for all vertices v of \mathbb{H} :

$$\alpha_1(v) \cap \alpha_2(v) = \alpha_0(v).$$

Note that 2-acyclicity of \mathbb{H} implies that any α -component $\mathbb{H}[\alpha; v]$ of \mathbb{H} , as an α -graph, is not only a weak but an induced substructure $\mathbb{H}[\alpha; v] \subseteq \mathbb{H}$: distinct vertices of $\mathbb{H}[\alpha; v]$ are by definition linked by an α -walk; an additional e -edge between them for $e \notin \alpha$ would therefore constitute a 2-cycle.

Observation 4.4. *Let \mathbb{H}_i, v_i for $i = 0, 1, 2$ be such that the free amalgam $\mathbb{H}, v := \mathbb{H}_1, v_1 \oplus_{\mathbb{H}_0, v_0} \mathbb{H}_2, v_2$ is well-defined, with embeddings $\rho_i: \mathbb{H}_0, v_0 \simeq \rho_i(\mathbb{H}_0), v_i \subseteq_w \mathbb{H}_i, v_i$ for $i = 1, 2$. If $\rho_i(\mathbb{H}_0) \subseteq_w \mathbb{H}_i$ is such that for $i = 1, 2$ every β -component of \mathbb{H}_i intersects the overlap region $\rho(\mathbb{H}_0) \subseteq_w \mathbb{H}_i$ in at most a single β -component of $\rho_i(\mathbb{H}_0) \simeq \mathbb{H}_0$, then the β -components in \mathbb{H} are (isomorphic to) β -components of just one of the \mathbb{H}_i or isomorphic to free amalgams of two β -components, one from each \mathbb{H}_i over a β -component of \mathbb{H}_0 ; in particular*

$$\mathbb{H}[\beta; v] \simeq \mathbb{H}_1[\beta; v_1] \oplus_{\mathbb{H}_0[\beta; v_0], v_0} \mathbb{H}_2[\beta; v_2].$$

The conditions in the observation are satisfied in particular for free amalgams of the form $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2]$, for $\mathbb{C}\mathbb{G}[\alpha_i] \subseteq_w \mathbb{C}\mathbb{G}$ over $\mathbb{C}\mathbb{G}[\alpha_0]$ for $\alpha_0 = \alpha_1 \cap \alpha_2$, if both $\mathbb{G}[\alpha_i]$ are 2-acyclic. It is also implied by 2-acyclicity of the constituent E-graphs \mathbb{H}_i for $i = 1, 2$ whenever these two E-graphs allow for a free amalgam over $\mathbb{H}_0, v_0 \simeq \mathbb{H}_i[\alpha_0; v_i], v_i$.

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 constituents \mathbb{H}_i in \mathbb{H} , then it is isomorphic to a β -component of that \mathbb{H}_i . Otherwise B must contain a vertex in the overlap of the constituent \mathbb{H}_i , and we consider w.l.o.g. the case of $v \in B$. For $i = 0, 1, 2$ let B_i be the vertex set of the β -component $\mathbb{H}_i[\beta; v_i]$ of v_i in \mathbb{H}_i , viewed as a subset of the vertex set V of the amalgam \mathbb{H} in terms of the natural embeddings. Clearly $B_0 \subseteq B_i \subseteq B$ for $i = 1, 2$. Under the assumptions of the observation, moreover, $B_0 = B_1 \cap B_2$ and this further implies that $B = B_1 \cup B_2$. It remains to argue that, as a β -graph, $\mathbb{H}[\beta; v]$ is isomorphic to the free amalgam of $\mathbb{H}_1[\beta; v_1]$ and $\mathbb{H}_2[\beta; v_2]$ over $\mathbb{H}_0[\beta; v_0]$. We first observe that this amalgam is well-defined, by means of the restrictions of the embeddings ρ_i to $\mathbb{H}_0[\beta; v_0] \subseteq_w \mathbb{H}_0$ (in particular condition $(*)$ from Definition 4.1 naturally restricts to $e \in \beta$). Now $\mathbb{H}_i \upharpoonright B_i$ as a β -graph is isomorphic to $\mathbb{H}_i[\beta; v_i]$, and by assumption, the embedded copies of these, for $i = 1, 2$, in \mathbb{H} overlap precisely in the embedded copy of $\mathbb{H}_0[\beta; v_0]$. \square

We turn to two specific forms of amalgams of Cayley graphs of generated subgroups: *amalgamation chains* and *amalgamation clusters* in § 4.1 and § 4.2, respectively. 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 $\mathbb{C}\mathbb{G}[\alpha_i]$ in the ‘tree-like’ pattern encountered in sufficiently acyclic \mathbb{G} .

4.1 Amalgamation chains

Definition 4.5. [amalgamation chain]

Let \mathbb{G} be an E-group, $N \geq 1$ and $\alpha_i \not\subseteq E$ for $1 \leq i \leq N$ with intersections $\alpha_{i,i+1} := \alpha_i \cap \alpha_{i+1}$ for $1 \leq i < N$. Consider the sequence of pointed E-graphs $(\mathbb{C}\mathbb{G}[\alpha_i], g_i)_{1 \leq i \leq N}$ and assume that the $g_i \in \mathbb{G}[\alpha_i] \subseteq_w \mathbb{C}\mathbb{G}$ 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* $\bigoplus_{i=1}^N (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ is the E-graph \mathbb{H} obtained as the result of simultaneous free amalgamation of disjoint copies of relational structures $\mathbb{H}_i \simeq \mathbb{C}\mathbb{G}[\alpha_i]$ via the shared embedded weak substructures $\mathbb{H}_{i,i+1} \simeq g_i\mathbb{C}\mathbb{G}[\alpha_{i,i+1}] \subseteq_w \mathbb{C}\mathbb{G}[\alpha_i]$ and $\mathbb{H}_{i,i+1} \simeq 1\mathbb{C}\mathbb{G}[\alpha_{i,i+1}] \subseteq_w \mathbb{C}\mathbb{G}[\alpha_{i+1}]$.

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 $1 \leq i < N$ are graph-theoretic separators along the chain.

Observation 4.6. *Whenever the amalgamation chain $\bigoplus (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ is defined, there is a unique homomorphism from $\bigoplus (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ to $\mathbb{C}\mathbb{G}$ that maps 1 in (the isomorphic copy of) $\mathbb{C}\mathbb{G}[\alpha_1]$ to $1 \in \mathbb{C}\mathbb{G}$. If the E-group \mathbb{G} is N -acyclic then this homomorphism is injective and $\bigoplus (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ is realised as a weak substructure of $\mathbb{C}\mathbb{G}$ (via an essentially canonical isomorphism).*

\mathbb{G} is N -acyclic if, and only if, $\bigoplus_i (\mathbb{C}\mathbb{G}[\alpha_i], g_i) \subseteq_w \mathbb{C}\mathbb{G}$ 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 $\bigoplus_{i=1}^n (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ that is not injectively mapped into $\mathbb{C}\mathbb{G}$ by the natural homomorphism. For an ℓ that is minimal with the property that $1 \leq k < k + \ell \leq N$ and that the subchain $\bigoplus_{i=k}^{\ell} (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ is not injectively embedded, its homomorphic image in $\mathbb{C}\mathbb{G}$ constitutes a coset cycle.

We are interested in the structure of β -components of amalgamation chains (in their rôle as partial unfoldings of coset cycles) towards compatibility guarantees in the inductive elimination of coset cycles over $\Gamma_k(\mathbb{G})$. For the simple situation of just a binary free amalgam $\mathbb{C}\mathbb{G}[\alpha_1] \oplus \mathbb{C}\mathbb{G}[\alpha_2]$, with overlap $\mathbb{C}\mathbb{G}[\alpha_0]$ for $\alpha_0 := \alpha_1 \cap \alpha_2$, the following is just a special case of Observation 4.4.

Lemma 4.7. *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 $\mathbb{C}\mathbb{G}[\beta \cap \alpha_1] \oplus \mathbb{C}\mathbb{G}[\beta \cap \alpha_2]$.*

Lemma 4.8. *Let $\alpha_i \subsetneq E$ for $1 \leq i \leq N$, $\beta \subseteq E$, 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 $\bigoplus_{i=s}^t (\mathbb{C}\mathbb{G}[\beta_i], h_i)$ for $\beta_i := \beta \cap \alpha_i$, some $1 \leq s \leq t \leq N$, and suitable choices of $h_i \in \mathbb{G}[\alpha_i]$.*

Corollary 4.9. *For any E-group \mathbb{G} , the functor that maps families of 2-acyclic $\mathbb{C}\mathbb{G}[\alpha_i]$ for $\alpha_i \subseteq E$ to free amalgamation chains of length up to N is conservative in the sense of Definition 2.13.*

Proof of Lemma 4.8. The claim follows by induction, through repeated application of Observation 4.4, based on the auxiliary claim that any β -component of $\bigoplus_{i=1}^n (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ for $n < N$ can intersect the $(\alpha_n \cap \alpha_{n+1})$ -component of the vertex representing $g_i \in \mathbb{C}\mathbb{G}[\alpha_i]$ in $\bigoplus_{i=1}^n (\mathbb{C}\mathbb{G}[\alpha_i], g_i)$ (i.e. the overlap in the next amalgamation step) in at most a single $(\beta \cap \alpha_n \cap \alpha_{n+1})$ -component (i.e. in a single β -component of that overlap). The first non-trivial instance, for the minimal n for which the β -component in question intersects the constituent $\mathbb{C}\mathbb{G}[\alpha_n]$ -copy, uses 2-acyclicity of that $\mathbb{G}[\alpha_n]$; the induction step from m to $m + 1$ similarly uses 2-acyclicity of the next $\mathbb{C}\mathbb{G}[\alpha_i]$ -copy in line, i.e. of $\mathbb{G}[\alpha_{n+1}]$. \square

4.2 Amalgamation clusters

While amalgamation chains reflect the free linear or path-like unfolding pattern of overlaps between cosets, amalgamation clusters reflect the overlap pattern of cosets w.r.t. their branching in individual elements. The following definition of free amalgamation clusters $\bigoplus_{i \in I} \mathbb{C}\mathbb{G}[\alpha_i]$ for a collection of generator sets $\alpha_i \subseteq E$ is similar in spirit to Definition 4.1 in capturing the idea of a minimal identification.

Definition 4.10. [free amalgamation cluster]

The *free amalgamation cluster* $\bigoplus_{i \in I} \mathbb{C}\mathbb{G}[\alpha_i]$ of Cayley graphs stemming from subgroups $\mathbb{G}[\alpha_i] \subseteq \mathbb{G}$ of the same E-group \mathbb{G} is obtained from the disjoint union of these $\mathbb{C}\mathbb{G}[\alpha_i]$ (als α_i -graphs) by identification of vertices from $\mathbb{C}\mathbb{G}[\alpha_i]$ and $\mathbb{C}\mathbb{G}[\alpha_{i'}]$ precisely if they stem from $\mathbb{G}[\alpha_i \cap \alpha_{i'}]$. We speak of a *free amalgamation cluster over* $\Gamma_k(\mathbb{C}\mathbb{G})$ if $\alpha_i \in \Gamma_k$ for all $i \in I$.

This definition does not immediately reduce to iterated application of the binary free amalgamation operation in the most general case. It does reduce to that, though, and with a result that is independent of the order and precedence of pairwise amalgamation steps, if the constituents are 2-acyclic, which will always be the case in further applications. If \mathbb{G} itself is 2-acyclic, then it embeds all free amalgamation clusters as weak substructures, and as induced substructures if it is 3-acyclic. This is another generalisation of Observation 4.2.

Observation 4.11. *If $\hat{\mathbb{G}} \succ_k \mathbb{G}$ is at least 3-acyclic w.r.t. Γ_k -generated cosets, i.e. admits no 2- or 3-cycles of Γ_k -generated cosets, then any free amalgamation cluster over $\Gamma_k(\mathbb{C}\mathbb{G})$ embeds as an induced substructure into $\mathbb{C}\hat{\mathbb{G}}$.⁵*

Proof. Existence of (an essentially unique) embedding of each constituent coset in the amalgamation cluster as a weak substructure follows from the identity $\Gamma_k(\hat{\mathbb{G}}) = \Gamma_k(\mathbb{G})$. That a free amalgamation cluster embeds as a weak substructure is guaranteed by 2-acyclicity of $\hat{\mathbb{G}}$ w.r.t. Γ_k -generated cosets: intersections of Γ_k -generated cosets in $\hat{\mathbb{G}}$ are the cosets generated by the common generators, as in free amalgams. That the embedding is onto an induced rather than just a weak substructure follows from 3-acyclicity of $\hat{\mathbb{G}}$ w.r.t. Γ_k -generated cosets: if images of two vertices from the same constituent $\mathbb{C}\mathbb{G}[\alpha]$ of the amalgamation cluster were related by an e -edge for $e \notin \alpha$ then this configuration would establish a 2-cycle of an α -generated and an $\{e\}$ -generated coset; if images of two vertices from the symmetric difference of the images of two distinct constituents $\mathbb{C}\mathbb{G}[\alpha_i]$ were related by an e -edge for $e \notin \alpha_1 \cup \alpha_2$, this would establish a 3-cycle formed by the α_i -generated cosets and an $\{e\}$ -generated coset. \square

Corollary 4.12. *Let \mathbb{G} admit a 3-acyclic $\hat{\mathbb{G}} \succ_k \mathbb{G}$.⁶ Then any free amalgamation cluster over $\Gamma_k(\mathbb{C}\mathbb{G})$ embeds into the Cayley graph $\mathbb{C}\hat{\mathbb{G}}$ of $\hat{\mathbb{G}}$ as an induced substructure.*

Again, we are interested in the consequences of this observation for connectivity phenomena within free amalgamation clusters over $\Gamma_k(\mathbb{C}\mathbb{G})$. In essence, if \mathbb{G} is compatible with all E-graphs in $\Gamma_k(\mathbb{C}\mathbb{G})$ (i.e. of Γ_k -generated subgroups), then internal connectivity patterns w.r.t. generator sets from Γ_k are strongly constrained by (in fact largely conform to) connectivity patterns in any 3-acyclic surrounding $\mathbb{C}\hat{\mathbb{G}}$, as provided by an embedding according to Corollary 4.12.

⁵As always, the embedding is unique up to arbitrary choice of the image of any one vertex of the connected E-graph to be embedded.

⁶We shall see in Section 5 that this is guaranteed if \mathbb{G} is compatible with the Cayley graphs of all its Γ_k -generated subgroups, which by Lemma 3.5 implies that \mathbb{G} itself does not admit 2- or 3-cycles of Γ_k -generated cosets; in particular all Γ_k -generated subgroups must be 3-acyclic.

Lemma 4.13. *Let \mathbb{G} admit a 3-acyclic $\hat{\mathbb{G}} \succ_k \mathbb{G}$.⁶ This implies the following for any free amalgamation cluster $\mathbb{H} = \bigoplus_{i \in I} \mathbb{C}\mathbb{G}[\alpha_i]$ with $\alpha_i \in \Gamma_k$ and $\beta, \gamma \in \Gamma_k$:*

- (i) *if \mathbb{H} is embedded into the 3-acyclic unfolding $\hat{\mathbb{G}} \succ_k \mathbb{G}$ according to Corollary 4.12, with vertex set $V \subseteq \hat{\mathbb{G}}$ for \mathbb{H} , then the vertex set $\beta(g) \subseteq V$ of the β -component of some g in \mathbb{H} is $V \cap g\hat{\mathbb{G}}[\beta]$.*
- (ii) *any β -connected component of \mathbb{H} (as an induced subgraph of \mathbb{H}) is isomorphic to a free amalgamation cluster of the form $\mathbb{H}_0 = \bigoplus_{i \in J} \mathbb{C}\mathbb{G}[\beta_i]$ for some $J \subseteq I$ and $\beta_i := \beta \cap \alpha_i$.*
- (iii) *the intersection of any β - and γ -components of \mathbb{H} is either empty or consists of a single $(\beta \cap \gamma)$ -component of \mathbb{H} , i.e. \mathbb{H} is 2-acyclic.*

Corollary 4.14. *The functor that maps families from $\Gamma_k(\mathbb{C}\mathbb{G})$ to free amalgamation clusters is conservative in the sense of Definition 2.13 in restriction to all E-groups \mathbb{G} that admit 3-acyclic unfoldings $\hat{\mathbb{G}} \succ_k \mathbb{G}$.⁶*

Proof of the lemma. We think of \mathbb{H} as embedded as $\mathbb{H} \subseteq \mathbb{C}\hat{\mathbb{G}}$ into the Cayley graph of a 3-acyclic E-group $\hat{\mathbb{G}}$ according to Corollary 4.12, rooted at $1 \in \hat{\mathbb{G}}$ say. Let $B \subseteq V$ be the vertex set of the β -component of \mathbb{H} in question. All α_i - and β_i -cosets in question can be equally regarded as cosets in $\hat{\mathbb{G}}$ as well as in \mathbb{G} , since $\mathbb{G} \preccurlyeq_k \hat{\mathbb{G}}$.

Assertion (i) is clear from 3-acyclicity of $\hat{\mathbb{G}}$. That $B \subseteq g\hat{\mathbb{G}}[\beta]$ is obvious even from $\mathbb{H} \subseteq_w \mathbb{C}\hat{\mathbb{G}}$. If this coset $g\hat{\mathbb{G}}[\beta]$ intersects two distinct constituents $\mathbb{G}[\alpha_i] = \mathbb{C}\mathbb{G}[\alpha_i]$ of \mathbb{H} , it must not form a 3-cycle together with these $\hat{\mathbb{G}}[\alpha_i]$; it follows that the β_i -cosets that make up the overlap between $g\hat{\mathbb{G}}[\beta]$ and the $\hat{\mathbb{G}}[\alpha_i]$ have to intersect in the intersection of the $\hat{\mathbb{G}}[\alpha_i]$ (and hence inside \mathbb{H}). So all intersections of $g\hat{\mathbb{G}}[\beta]$ with constituents $\mathbb{G}[\alpha_i]$ of \mathbb{H} belong to the same β -component of \mathbb{H} .

For (ii) we may use (i) to represent the vertex set $B \subseteq V$ of the β -component as $V \cap g\hat{\mathbb{G}}[\beta] = V \cap g\mathbb{G}[\beta] \subseteq \hat{\mathbb{G}}$. Any constituent $\mathbb{C}\mathbb{G}[\alpha_i]$ of \mathbb{H} that intersects B in some vertex $g_i \in B \cap \mathbb{G}[\alpha_i]$ contributes the full β_i -coset $g_i\mathbb{G}[\beta_i]$ to B . Moreover, each constituent $\mathbb{C}\mathbb{G}[\alpha_i]$ intersects B in at most one such β_i -coset, by 2-acyclicity. We may therefore choose vertices $g_j \in B \cap \hat{\mathbb{G}}[\alpha_j]$ for all $j \in J \subseteq I$, where J is the set of those indices $i \in I$ for which B intersects $\hat{\mathbb{G}}[\alpha_i]$, so that $B = \bigcup_{j \in J} g_j\mathbb{G}[\beta_j]$. For $j \in J$, let $\alpha_{0j} \subseteq \alpha_j$ be the \subseteq -minimal set of generators for which $g_j\mathbb{G}[\beta_j]$ intersects $\mathbb{G}[\alpha_{0j}] \subseteq \mathbb{G}[\alpha_j]$. (Existence and uniqueness of such minimal subsets follow from 3-acyclicity of the $\mathbb{G}[\alpha_j]$.) So the α_{0j} -cosets at $1 \in \hat{\mathbb{G}}$ all contain vertices from B , i.e. they all intersect one and the same β -coset of $\hat{\mathbb{G}}$. Let $h_{0j} \in \mathbb{G}[\alpha_{0j}] \cap B$ for $j \in J$. Then $B = \bigcup_{j \in J} g_j\mathbb{G}[\beta_j] = \bigcup_{j \in J} h_{0j}\mathbb{G}[\beta_j]$. By 3-acyclicity of $\hat{\mathbb{G}}$, the respective β_j -cosets of any two such h_{0j} must intersect within B : for $j \neq j' \in J$, $\hat{\mathbb{G}}[\alpha_{0j}], \hat{\mathbb{G}}[\alpha_{0j'}]$ and $g\hat{\mathbb{G}}[\beta]$ must not form a 3-cycle in $\hat{\mathbb{G}}$, whence $\hat{\mathbb{G}}[\alpha_{0j}] \cap \hat{\mathbb{G}}[\alpha_{0j'}]$ must intersect in $g\hat{\mathbb{G}}[\beta]$, and as the intersection is contained in V , it meets B . It follows that we may choose $h_{0j} = h_{0j'}$, and that, by minimality, $\alpha_{0j} = \alpha_{0j'}$. As this applies to any pair of indices $j, j' \in J$, all α_{0j} are identical, and hence subsets of $\bigcap_{j \in J} \alpha_j$ for $j \in J$. Assuming w.l.o.g. that

$0 \in J \subseteq I$, we can use $\alpha_{0j} = \alpha_{00} \subseteq \alpha_j$ for all $j \in J$, and correspondingly choose all $h_{0j} \in \mathbb{G}[\alpha_{00}]$. By 3-acyclicity of $\hat{\mathbb{G}}$, all these h_{0j} are elements of a common β_{00} -coset for $\beta_{00} := \beta \cap \alpha_{00}$. With these h_{0j} , it follows that $B = \bigcup_{j \in J} h_{0j} \mathbb{G}[\beta_j]$ is an embedded free amalgam of copies of the $\mathbb{C}\mathbb{G}[\beta_j]$ for $j \in J$, embedded as an induced subgraph in $\mathbb{C}\hat{\mathbb{G}}$ and hence in \mathbb{H} as claimed.

Finally, (iii) is an obvious consequence of (ii). The vertex set of a non-trivial intersection between a β -component and a γ -component is a connected component w.r.t. $\beta \cap \gamma$. \square

The following slight generalisation will be useful for later constructions.

Lemma 4.15. *Let \mathbb{G} admit a 3-acyclic $\hat{\mathbb{G}} \succ_k \mathbb{G}$.⁶ Then all free amalgams of the form $\mathbb{H}, g \oplus \mathbb{C}\mathbb{G}[\beta]$ between a free amalgamation cluster $\mathbb{H} = \bigoplus_{i \in I} \mathbb{C}\mathbb{G}[\alpha_i]$ over $\Gamma_k(\mathbb{C}\mathbb{G})$ and $\mathbb{C}\mathbb{G}[\beta]$ from $\Gamma_k(\mathbb{C}\mathbb{G})$ are well defined and embed into $\mathbb{C}\hat{\mathbb{G}}$ as induced substructures: $\mathbb{H}, g \oplus \mathbb{C}\mathbb{G}[\beta] \subseteq \mathbb{C}\hat{\mathbb{G}}$. Moreover, for $\gamma \in \Gamma_k$, any γ -component of free amalgams $\mathbb{H}, v \oplus \mathbb{C}\mathbb{G}[\beta]$ is isomorphic to $\mathbb{C}\mathbb{G}[\gamma]$ or to a free amalgam of the form $\bigoplus_{i \in J} \mathbb{C}\mathbb{G}[\gamma_i], g' \oplus \mathbb{C}\mathbb{G}[\gamma']$ for $J \subseteq I$, $\gamma_i = \gamma \cap \alpha_i$ and $\gamma' \subseteq \gamma \cap \beta$ (possibly empty).*

Proof. The claims about existence and embeddability follow directly from the previous lemma. For the claim about γ -components of $\mathbb{H}, g \oplus \mathbb{C}\mathbb{G}[\beta]$, we may use 2-acyclicity of \mathbb{H} (by (iii) from the previous lemma) and of $\mathbb{C}\mathbb{G}[\beta]$ to apply Observation 4.4. \square

Corollary 4.16. *The functor that maps families from $\Gamma_k(\mathbb{C}\mathbb{G})$ to free amalgams between free amalgamation clusters over $\Gamma_k(\mathbb{C}\mathbb{G})$ and one individual $\mathbb{C}\mathbb{G}[\alpha]$ from $\Gamma_k(\mathbb{C}\mathbb{G})$ is conservative in the sense of Definition 2.13 in restriction to all E-groups \mathbb{G} that admit 3-acyclic unfoldings $\hat{\mathbb{G}} \succ_k \mathbb{G}$.⁶*

5 Construction of N-acyclic E-groups

We may construct N -acyclic E-groups based on the general considerations in Proposition 2.14. The functor F , which collects the obstructions at a given level Γ_m into an E-graph \mathbb{H} that eliminates these obstructions at level Γ_{m+1} , just builds amalgamation chains of appropriate length over $\Gamma_m(\mathbb{C}\mathbb{G})$. Its conservative nature was established in Section 4.1, Lemma 4.8.

In essence we find, for fixed $N \geq 2$, a sequence of E-groups \mathbb{G}_m such that

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

i.e. \mathbb{G}_m admits no coset cycles of length up to N with constituents $\Gamma_m(\mathbb{C}\mathbb{G}_m)$ (α -cosets for $|\alpha| < m$). In order to preserve the acyclicity property at lower levels in the induction step, we pass from \mathbb{G}_m to $\mathbb{G}_m \preccurlyeq_m \mathbb{G}_{m+1}$ obtain \mathbb{G}_{m+1} as $\mathbb{G}_{m+1} := \text{sym}(\mathbb{H}_m)$, with \mathbb{H}_m the disjoint union of $\mathbb{C}\mathbb{G}_m$ and $F(\Gamma_m(\mathbb{C}\mathbb{G}_m))$.

The following determines the appropriate length bound, which generalises Lemma 3.5.

Lemma 5.1. *For $n \geq 1$, 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 $\mathbb{G} \succ \text{sym}(\mathbb{H})$ 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 $\text{sym}(\mathbb{H})$ 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 $n = N - 2$. 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 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 $[\prod_{i=0}^{n+1} w_i]_{\mathbb{H}} \in \text{sym}(\mathbb{H})$, 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 connectivity and separation conditions 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 $[\prod_{i=0}^{n+1} w_i]_{\mathbb{G}} = 1$. \square

Proposition 5.2. *Let $n \geq 1$ and let F be the functor that maps an E -group \mathbb{G} to the collection of all free amalgamation chains of length up to n over constituents $\mathbb{C}\mathbb{G}[\alpha] \subseteq_w \mathbb{C}\mathbb{G}$ for $\alpha \subseteq E$. For fixed $1 \leq k \leq |E| + 1$ consider an E -group \mathbb{G} that is compatible with the E -graphs in $F(\Gamma_k(\mathbb{G}))$.⁷ Then \mathbb{G} admits an unfolding $\hat{\mathbb{G}} \succ_k \mathbb{G}$ that is N -acyclic for $N = n + 2$.*

Proof. This is a direct application of Proposition 2.14 to the functor F , which is conservative according to Lemma 4.8. \square

Recall from Definition 2.15 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 Proposition 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

⁷Note that this is vacuously true at level $k = 1$.

\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 proposition.

Theorem 5.3. *For every finite E-group \mathbb{G} and $N \geq 2$ 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 $\mathbb{G} := \text{sym}(\mathbb{H})$ 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{aligned} [\]_{\mathbb{G}}: 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{aligned}$$

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 $\mathbb{C}\mathbb{G}$ 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 $\mathbb{G} \succcurlyeq \text{sym}(\mathbb{I})$.*

Remark 6.2. *The restriction to E-groups that are compatible with \mathbb{I} , $\mathbb{G} \succcurlyeq \text{sym}(\mathbb{I})$, 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 $\pi_e \in \text{sym}(\mathbb{I})$. As \mathbb{I} will typically not be a complete E-graph, the Cayley graph of $\mathbb{C}\mathbb{G}$ does not map homomorphically to \mathbb{I} . But α -walks from $s \in \mathbb{I}$ do have unique lifts to α -walks from any $g \in \mathbb{C}\mathbb{G}$. Compatibility says that lifts at the same $g \in \mathbb{C}\mathbb{G}$ of different walks from $s \in S$ can only meet in $\mathbb{C}\mathbb{G}$ 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 notion 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

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

with $s_i \in S$ and $e_i \in E$ where $(s_i, s_{i+1}) \in R_{e_{i+1}}$ for $0 \leq i < n$ 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 $\mathbb{I} \upharpoonright \alpha$), 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} .

In the following, $\alpha^*[\mathbb{I}]$ -walks will also be addressed as $\mathbb{I}[\alpha]$ -walks, especially if the restriction of the available generators to $\alpha \subseteq E$ seems more important than the actual labelling.

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, but keep in mind that, e.g. $\mathbb{I}[\alpha]$ -walks inside such $\mathbb{I}[\alpha]$ -components implicitly refer to a fixed choice of an anchor point that determines which edges are available.

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

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

Note that all α -walks in $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$, 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 $\mathbb{C}\mathbb{G}$.

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

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 $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha, s; g] \subseteq_w \mathbb{C}\mathbb{G}$ of g .

We drop mention of g for cosets at $g = 1$, writing e.g. just $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha, s] \subseteq_w \mathbb{C}\mathbb{G}$ for $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha, s; 1] \subseteq_w \mathbb{C}\mathbb{G}$. 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, $n \geq 2$. An \mathbb{I} -coset cycle of length $n \geq 2$ in \mathbb{G} or $\mathbb{C}\mathbb{G}$ is a cyclically indexed sequence $(\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i; g_i], g_i)_{i \in \mathbb{Z}_n}$ 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} ,⁸
- (ii) (separation) $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_{i,i-1}, s_i; g_i] \cap \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_{i,i+1}, s_{i+1}; g_{i+1}] = \emptyset$,

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

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 subsumed 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 ,

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

⁸Equivalently, $g_{i+1} \in \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i; g_i]$, or $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i; g_i] = \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_{i+1}; g_{i+1}]$

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 idea is to unfold potential \mathbb{I} -coset cycles with constituents $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i, g_i] \subseteq_w \mathbb{C}\mathbb{G}[\alpha_i]$ through the unfolding of the surrounding plain cosets $\mathbb{C}\mathbb{G}[\alpha_i]$ 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 new challenge is to safeguard this unfolding and the corresponding compatibility arguments against unwanted shortcuts by α -walks that do not correspond to $\mathbb{I}[\alpha]$ -walks.

It is to this end that we consider the relationship between $\mathbb{I}[\alpha]$ -components and the halo around them that is formed by overlapping (unconstrained) α' -connected components for $\alpha' \subsetneq \alpha$. The \mathbb{I} -*skeletons* to be defined below play the rôle of $\mathbb{I}[\alpha]$ -connected components that unfold all those α -links that are available according to \mathbb{I} , and none else. These arise naturally as $\mathbb{I}[\alpha]$ -components of Cayley graphs $\mathbb{C}\mathbb{G}$ provided \mathbb{G} is compatible with \mathbb{I} , as we assumed with Proviso 6.1. But we also consider individual such \mathbb{I} -skeletons as scaffoldings for the construction of further E-graphs, which are to be fed into the unfolding process that successively unclutters the embedded $\mathbb{I}[\alpha]$ -components and unfolds short \mathbb{I} -coset cycles for increasing numbers of generators.

Definition 6.10. [\mathbb{I} -skeleton]

For $\alpha \subseteq E$ an $\mathbb{I}[\alpha, s]$ -*skeleton* is a connected α -graph \mathbb{H} that admits a surjective homomorphism $h: \mathbb{H} \rightarrow \mathbb{I}[\alpha, s]$ 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} . Across all $s \in S$ we speak of just $\mathbb{I}[\alpha]$ -skeletons.

When speaking of an $\mathbb{I}[\alpha]$ -skeleton, it is important to keep in mind that it is an $\mathbb{I}[\alpha, s]$ -skeleton for some $s \in S$.

Example 6.11. Every $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha, s; g] \subseteq_w \mathbb{C}\mathbb{G}[\alpha] \subseteq_w \mathbb{C}\mathbb{G}$ 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). We refer to such instances as *embedded \mathbb{I} -skeletons*.

The lifting property in the definition guarantees that every $\mathbb{I}[\alpha]$ -walk from s' in $\mathbb{I}[\alpha, s]$ has a unique lift to an α^* - \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 $\mathbb{I} \upharpoonright \alpha$).

Proviso 6.12. *In order not to overburden notation, we often leave implicit the S -marking of the elements of an $\mathbb{I}[\alpha, s]$ -skeleton \mathbb{H} that comes with a given homomorphism $h: \mathbb{H} \rightarrow \mathbb{I}[\alpha, s]$. Speaking of an $\mathbb{I}[\alpha]$ -walk in \mathbb{H} , for instance, we only admit walks labelled by $w \in \alpha^*[\mathbb{I}, s, t]$ from v to v' in \mathbb{H} if these are marked accordingly, i.e. for $h(v) = s$, which then implies $h(v') = t$.*

Definition 6.13. [freeness]

The embedded $\mathbb{I}[\alpha]$ -skeleton $\mathbb{H} := \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha, s; g] \subseteq_w \mathbb{C}\mathbb{G}[\alpha] \subseteq_w \mathbb{C}\mathbb{G}$ is *free* in $\mathbb{C}\mathbb{G}[\alpha]$ or in $\mathbb{C}\mathbb{G}$ if for any two elements $v_1, v_2 \in \mathbb{H}$, any cosets $v_1\mathbb{G}[\alpha_1]$ and

$v_2\mathbb{G}[\alpha_2]$, for $\alpha_i \subsetneq \alpha$, overlap in the surrounding $\mathbb{C}\mathbb{G}$, if, and *only if*, the α_i -connected components of these elements v_i overlap within \mathbb{H} :

$$\alpha_1(v_1) \cap \alpha_2(v_2) = \emptyset \text{ in } \mathbb{H} \implies v_1\mathbb{G}[\alpha_1] \cap v_2\mathbb{G}[\alpha_2] = \emptyset \text{ in } \mathbb{C}\mathbb{G}[\alpha].$$

The E-group \mathbb{G} is *free over* \mathbb{I} if every embedded $\mathbb{I}[\alpha]$ -skeleton for $\alpha \subseteq E$ is free in $\mathbb{C}\mathbb{G}$.

Already the special case of $\alpha' := \alpha_1 = \alpha_2 \subsetneq \alpha$ gives an indication of the strength of this freeness condition: the freeness requirement here says that α' -reachability inside the embedded $\mathbb{I}[\alpha]$ -skeleton agrees with α' -reachability in the surrounding $\mathbb{C}\mathbb{G}$. The following is an analogue of Lemma 3.6 for freeness.

Lemma 6.14. *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 $h: \hat{\mathbb{G}} \rightarrow \mathbb{G}$ 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 $\mathbb{C}\mathbb{G}$. Assume that an $\mathbb{I}[\alpha, s]$ -skeleton $\hat{\mathbb{H}} \subseteq_w \mathbb{C}\hat{\mathbb{G}}$ violates the freeness condition through some $\hat{g} \in v_1\hat{\mathbb{G}}[\alpha_1] \cap v_2\hat{\mathbb{G}}[\alpha_2] \neq \emptyset$ where $\mathbb{C}\hat{\mathbb{G}}[\mathbb{I}, \alpha_1, s_1; v_1] \cap \mathbb{C}\hat{\mathbb{G}}[\mathbb{I}, \alpha_2, s_2; v_2] = \emptyset$. By assumption h is injective in restriction to each one of the two cosets $v_i\hat{\mathbb{G}}[\alpha_i]$ 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 $\hat{g}\hat{\mathbb{G}}[\alpha_0]$ 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 $\mathbb{C}\hat{\mathbb{G}}[\alpha]$ would be isomorphically mapped onto a violation in $\mathbb{G}[\alpha]$. \square

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 6.17 below, consider $\alpha, \alpha_1, \alpha_2$ 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 6.15. *For $\alpha_1, \alpha_2 \subsetneq \alpha$ consider vertices v_1, v_2 of an embedded $\mathbb{I}[\alpha]$ -skeleton $\mathbb{H} := \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha, s_1; v_1] = \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha, s_2; v_2] \subseteq_w \mathbb{C}\mathbb{G}[\alpha]$ that is free in $\mathbb{C}\mathbb{G}$. If $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_1, s; v_1] \cap \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_2, s'; v_2] = \emptyset$ (separation as in \mathbb{I} -coset cycles) then also $v_1\mathbb{G}[\alpha_1] \cap v_2\mathbb{G}[\alpha_2] = \emptyset$ (separation as in plain coset cycles).*

Finally, the condition on overlaps in Definition 6.13 takes an especially neat and concrete form if \mathbb{G} is 2-acyclic.

Observation 6.16. *Let the embedded $\mathbb{I}[\alpha]$ -skeleton $\mathbb{H} := \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha, s; g] \subseteq_w \mathbb{C}\mathbb{G}$ be free in $\mathbb{C}\mathbb{G}[\alpha]$ and let $\mathbb{G}[\alpha]$ be 2-acyclic. Then two cosets $v_1\mathbb{G}[\alpha_1]$ and $v_2\mathbb{G}[\alpha_2]$ for $v_i \in \mathbb{H}$ and $\alpha_i \subsetneq \alpha$ are either disjoint or their overlap is a single coset of the form $v_0\mathbb{G}[\alpha_0]$ for $\alpha_0 := \alpha_1 \cap \alpha_2$ and an element $v_0 \in \mathbb{H}$ that is α_i -reachable from v_i inside the skeleton, for $i = 1, 2$.*

Proof. Due to 2-acyclicity a non-empty intersection of α_i -cosets in $\mathbb{C}\mathbb{G}[\alpha]$ is a single α_0 -coset. By freeness it must intersect the embedded $\mathbb{I}[\alpha]$ -skeleton in an element v_0 that is α_i -reachable from v_i in the skeleton. \square

Lemma 6.17. *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 N -acyclic 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 a plain coset cycle if the links $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i; g_i]$ are extended to the surrounding $g_i\mathbb{C}\mathbb{G}[\alpha_i]$. The crux is the separation condition: separation w.r.t. \mathbb{I} -reachability in links $\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i; g_i] \subseteq_w g_i\mathbb{C}\mathbb{G}[\alpha_i]$ 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 6.15. \square

7 Small coset amalgams over \mathbb{I} -skeletons

We think of a fixed constraint graph \mathbb{I} and its generated subgraphs $\mathbb{I}[\alpha]$, and only consider E-groups \mathbb{G} that are compatible with \mathbb{I} according to Proviso 6.1.

For *small coset amalgams* we want to amalgamate constituent $\mathbb{C}\mathbb{G}[\alpha']$ -copies for $\alpha' \subsetneq \alpha$ along an $\mathbb{I}[\alpha]$ -skeleton $\mathbb{H}[\alpha]$ to form an E-graph, in as free a manner as possible. In every vertex v of the skeleton, we mean to amalgamate v -tagged copies $v\mathbb{C}\mathbb{G}[\alpha']$ of $\mathbb{C}\mathbb{G}[\alpha']$ with minimal identifications as are enforced by the requirements for E-graphs. Before looking at embedded realisations of such amalgams in a surrounding $\mathbb{C}\mathbb{G}[\alpha]$ (with appropriate acyclicity conditions) we define these amalgams in isolation (with only intrinsic acyclicity conditions for the skeleton $\mathbb{H}[\alpha]$ in relation to the constituents $\mathbb{C}\mathbb{G}[\alpha']$ for $\alpha' \subsetneq \alpha$). The prerequisite (iii) in the following definition corresponds to the freeness condition for embedded \mathbb{I} -skeletons of Definition 6.13, here expressed for α' -components $\mathbb{H}[\alpha'; v] \subseteq_w \mathbb{H}[\alpha]$ for $\alpha' \subsetneq \alpha$.

Definition 7.1. [small coset amalgam]

An $\mathbb{I}[\alpha]$ -skeleton $\mathbb{H} := \mathbb{H}[\alpha]$ and the family of subgroups $\mathbb{G}[\alpha']$ for $\alpha' \subsetneq \alpha$ admit a *small coset amalgam* if

- (i) $\mathbb{H}[\alpha]$ is 2-acyclic;
- (ii) the $\mathbb{G}[\alpha']$ are 2-acyclic;
- (iii) all α' -components of $\mathbb{H}[\alpha]$ embed isomorphically (as α' -graphs) into $\mathbb{C}\mathbb{G}[\alpha']$ in a free manner, i.e. there is an embedding $\rho: \mathbb{H}[\alpha'; v] \rightarrow \mathbb{C}\mathbb{G}[\alpha']$ such that for $\alpha'_1, \alpha'_2 \subsetneq \alpha'$:

$$\alpha'_1(v_1) \cap \alpha'_2(v_2) = \emptyset \text{ in } \mathbb{H}[\alpha'; v] \implies \alpha'_1(\rho(v_1)) \cap \alpha'_2(\rho(v_2)) = \emptyset \text{ in } \mathbb{C}\mathbb{G}[\alpha'].$$

Under these conditions the *small coset amalgam*

$$\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha) := \bigoplus_{v \in V, \alpha' \subsetneq \alpha} v\mathbb{C}\mathbb{G}[\alpha'] / \approx,$$




Figure 2: Pattern for identifications in $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$, see (\dagger) in Definition 7.1: for $i = 1, 2$, $g_{i0} = [w_{i0}]_{\mathbb{G}}$ links v_0 to v_i in their α'_i -component of \mathbb{H} .




Figure 3: Pattern for two consecutive identifications.

is defined as the quotient of the disjoint union of v -tagged constituents $v\mathbb{C}\mathbb{G}[\alpha']$ w.r.t. \approx , the equivalence relation induced by the following identifications, cf. Figure 2 and recall Proviso 6.12 for walks in $\mathbb{H}[\alpha]$:

- (\dagger) $v_1g_1 \in v_1\mathbb{C}\mathbb{G}[\alpha'_1]$ is identified with $v_2g_2 \in v_2\mathbb{C}\mathbb{G}[\alpha'_2]$ if for some v_0 in \mathbb{H} and $g_0 \in \mathbb{G}[\alpha'_0]$ for $\alpha'_0 := \alpha'_1 \cap \alpha'_2$, v_0 is reachable from v_i inside \mathbb{H} by an $\mathbb{I}[\alpha'_i]$ -walk labelled w_{i0} such that $v_0 = v_i[w_{i0}]_{\mathbb{H}}$ and $g_{i0}g_0 := [w_{i0}]_{\mathbb{G}}g_0 = g_i$ in $\mathbb{G}[\alpha'_i]$ for $i = 1, 2$.

Pre-condition (iii) in the definition is essential for the analysis of \approx for identification of elements according to (\dagger) . Let \sim denote the relation on the disjoint union of the v -tagged copies of constituent $v\mathbb{C}\mathbb{G}[\alpha']$: $v_1g_1 \sim v_2g_2$ if (\dagger) , where we write $v_i g_i$ instead of (v_i, g_i) for the element g_i in the v_i -tagged copy of $\mathbb{G}[\alpha'_i]$ that forms the constituent $v_i\mathbb{C}\mathbb{G}[\alpha'_i]$. Then \approx is the transitive closure of \sim . But condition (iii) entails that \sim itself is transitive so that (\dagger) directly describes all identifications in $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$. Transitivity of \sim is shown as follows, cf. Figure 3. Let $v_1g_1 \sim v_2g_2 \sim v_3g_3$ where $g_i \in \mathbb{G}[\alpha'_i]$ is an element of the constituent $v_i\mathbb{C}\mathbb{G}[\alpha'_i]$, for $\alpha'_i \in \Gamma_\alpha$, i.e. $\alpha'_i \subsetneq \alpha$. By the definition of \sim in (\dagger) this implies the existence of intermediaries $v_{i,j}g_{i,j}$ for $(i, j) = (1, 2), (2, 3)$ according to (\dagger) . I.e. $v_{i,j}$ is in the intersection of the α'_i -component of v_i and the α'_j -component

of v_j in \mathbb{H} and $g_{i,j} \in \mathbb{G}[\alpha_{i,j}]$ with $\alpha_{i,j} := \alpha'_i \cap \alpha'_j$ are such that (cf. Figure 3)

- $v_{1,2}g_{1,2}$ is identified with v_1g_1 in the $\mathbb{CG}[\alpha'_1]$ -copy at v_1 and with v_2g_2 in the $\mathbb{CG}[\alpha'_2]$ -copy at v_2 ;
- $v_{2,3}g_{2,3}$ is identified with v_2g_2 in the $\mathbb{CG}[\alpha'_2]$ -copy at v_2 and with v_3g_3 in the $\mathbb{CG}[\alpha'_3]$ -copy at v_3 .

We look at the identifications that occur in the central $\mathbb{CG}[\alpha'_2]$ -copy. That the α'_2 -component of \mathbb{H} containing $v_{1,2}$, v_2 and $v_{2,3}$ satisfies condition (iii) means that at least one of the following must hold

- (1) not $\alpha_{1,2} \subsetneq \alpha'_2$, which implies $\alpha'_2 \subseteq \alpha'_1$, or
- (2) not $\alpha_{2,3} \subsetneq \alpha'_2$, which implies $\alpha'_2 \subseteq \alpha'_3$, 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_1g_1 with v_3g_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. So the crucial assumption (iii) implies that \sim is itself transitive, hence equal to \approx .

Moreover, the isomorphic embeddings of α' -components of $\mathbb{H}[\alpha]$ into $\mathbb{CG}[\alpha']$ for $\alpha' \subsetneq \alpha$, which are required in (iii) of Definition 7.1, rule out that \sim induces any identifications within the skeleton $\mathbb{H}[\alpha]$: $v_11 \sim v_21$ can only be mediated by some v_0g_0 such that $g_0 = g_{10}^{-1} = g_{20}^{-1} \in \mathbb{G}[\alpha'_0]$ for $\alpha_0 = \alpha'_1 \cap \alpha'_2$. By condition (iii), once more, the v_i for $i = 0, 1, 2$ are in the same α'_0 -component of $\mathbb{H}[\alpha]$ which implies $v_1 = v_2$. It follows that $\mathbb{H}[\alpha] \subseteq_w \mathbf{CE}(\mathbb{H}, \mathbb{G}, \alpha)$. We summarise these findings in the following remark.

Remark 7.2. $\hat{\mathbb{H}} := \mathbf{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ as in Definition 7.1 embeds the $\mathbb{I}[\alpha]$ -skeleton $\mathbb{H}[\alpha]$ as a weak substructure, $\mathbb{H}[\alpha] \subseteq_w \mathbf{CE}(\mathbb{H}, \mathbb{G}, \alpha)$, and the equivalence relation \approx identifies elements from different constituents $v_i\mathbb{CG}[\alpha'_i]$ precisely if they are related as in (†).

Constituents $v\mathbb{CG}[\alpha']$ and $v'\mathbb{CG}[\alpha']$ coincide whenever v and v' are linked by an $\mathbb{I}[\alpha']$ -walk in $\mathbb{H}[\alpha]$ (α' -connected in $\mathbb{H}[\alpha]$).

We now look at strong *freeness* criteria for the manner in which $\mathbb{I}[\alpha]$ -skeletons and small coset amalgams embed into $\mathbb{CG}[\alpha]$. The basic freeness condition, which we just repeat in context as (i) here, is as in Definition 6.13. It also occurred as an intrinsic condition for the constituents $\mathbb{CG}[\alpha']$ for $\alpha' \subsetneq \alpha$ in (iii) of Definition 7.1 above.

Definition 7.3. [freeness and bridge-freeness]

An embedded $\mathbb{I}[\alpha]$ -skeleton $\mathbb{H}[\alpha] \subseteq_w \mathbb{CG}[\alpha] \subseteq \mathbb{CG}$ is

- (i) *free* if, for all $\alpha_1, \alpha_2 \subsetneq \alpha$ and v_1, v_2 from $\mathbb{H}[\alpha]$:

$$\alpha_1(v_1) \cap \alpha_2(v_2) = \emptyset \text{ in } \mathbb{H}[\alpha] \implies \alpha_1(v_1) \cap \alpha_2(v_2) = \emptyset \text{ in } \mathbb{CG}[\alpha].$$




Figure 4: Bridge-freeness.

- (ii) *bridge-free* if it is free and $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha) \subseteq_w \mathbb{CG}[\alpha]$ is well-defined as a weak substructure of $\mathbb{CG}[\alpha]$ over the embedded skeleton $\mathbb{H}[\alpha] \subseteq_w \mathbb{CG}[\alpha]$, and if for all $\beta \subsetneq \alpha$ and $v_1 g_1, v_2 g_2 \in \text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$: if $v_1 g_1$ and $v_2 g_2$ are β -connected in $\mathbb{G}[\alpha]$ (i.e. if $v_1 g_1 \mathbb{G}[\beta] = v_2 g_2 \mathbb{G}[\beta]$), then they are β -connected inside $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ in the following simple form (cf. Figure 4):

- (*) for some $\alpha_1, \alpha_2 \subsetneq \alpha$ such that $g_i \in \mathbb{G}[\alpha_i]$ for $i = 1, 2$ there is some $v_0 g_0 \in \text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$, where $g_0 \in \mathbb{G}[\alpha_0]$ for $\alpha_0 := \alpha_1 \cap \alpha_2$ and such that $v_i g_i \in v_0 g_0 \mathbb{G}[\beta_i]$ for $i = 1, 2$ where $\beta_i := \beta \cap \alpha_i$.

In essence, the connectivity condition in (ii) implies that β -connectivity inside the embedded $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha) \subseteq_w \mathbb{CG}[\alpha]$ coincides with β -connectivity in the surrounding $\mathbb{G}[\alpha] \subseteq \mathbb{G}$.

Provided that $\mathbb{G}[\alpha]$ and $\mathbb{H}[\alpha] \subseteq_w \mathbb{CG}[\alpha]$ are 2-acyclic, condition (i) implies that $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ embeds isomorphically as a weak substructure

$$\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha) \subseteq_w \mathbb{CG}[\alpha]$$

over the embedded $\mathbb{H}[\alpha]$. In this case, condition (ii) then additionally implies, by application to singleton $\beta = \{e\} \subseteq \alpha$, an induced substructure relationship

$$\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha) \subseteq \mathbb{CG}[\alpha].$$

We also note that, for 2-acyclic $\mathbb{G}[\alpha]$, the freeness of $\mathbb{H}[\alpha'] \subseteq_w \mathbb{CG}[\alpha']$ for $\alpha' \subsetneq \alpha$ implies that $\mathbb{H}[\alpha]$ is 2-acyclic (condition (i) of Definition 7.1). For this consider vertices v_1, v_2 of $\mathbb{H}[\alpha]$ that are α_i -connected in $\mathbb{H}[\alpha]$ for $\alpha_1, \alpha_2 \subsetneq \alpha$. If v_1 and v_2 were not part of the same α_0 -component of $\mathbb{H}[\alpha]$, for $\alpha_0 := \alpha_1 \cap \alpha_2$, then by freeness of $\mathbb{H}[\alpha_1] \subseteq_w \mathbb{CG}[\alpha_1]$, $v_1 \mathbb{CG}[\alpha_0] \cap v_2 \mathbb{CG}[\alpha_0] = \emptyset$; and this would contradict 2-acyclicity of $\mathbb{CG}[\alpha]$ which implies that v_1 and v_2 are α_0 -connected in $\mathbb{G}[\alpha]$. We state these observations for further reference.

Observation 7.4. *If $\mathbb{H}[\alpha'] \subseteq_w \mathbb{CG}[\alpha']$ is free for all $\alpha' \subsetneq \alpha$ and $\mathbb{G}[\alpha]$ is 2-acyclic, then $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ is well-defined.*

If $\mathbb{H}[\alpha] \subseteq_w \mathbb{C}\mathbb{G}[\alpha]$ itself is free, then $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ embeds as a weak substructure $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha) \subseteq_w \mathbb{C}\mathbb{G}[\alpha]$; if $\mathbb{H}[\alpha] \subseteq_w \mathbb{C}\mathbb{G}[\alpha]$ is even bridge-free, then $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha) \subseteq \mathbb{C}\mathbb{G}[\alpha]$ is embedded as an induced substructure.

We say that a set α of generators *supports* some vertex over some embedded skeleton if that vertex is linked to the skeleton by an α -walk. We extend this notion to sets of vertices, and consider minimal supports as in the following definition.

Definition 7.5. [support and unique minimal support]

If the small coset amalgam $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ is defined, and for a vertex w of $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ or for a subset B of $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$, or a subset $B \subseteq \mathbb{G}[\alpha]$ in case $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha) \subseteq_w \mathbb{C}\mathbb{G}[\alpha]$:

- (i) a $\mathbb{C}\mathbb{G}[\alpha']$ -constituent $v\mathbb{C}\mathbb{G}[\alpha']$ of $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ for $\alpha' \subsetneq \alpha$ *supports* a vertex w in $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ if $w \in v\mathbb{C}\mathbb{G}[\alpha']$; in this case we also say that the generator set α' provides support for w ; similarly for a set B of vertices, the requirement is that some element of B be supported in this manner.
- (ii) the set of generators $\alpha' \subsetneq \alpha$ provides *minimal support* for a vertex w in $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ if some $\mathbb{C}\mathbb{G}[\alpha']$ -constituent of $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ supports w , and if α' is a subset of every such generator set; similarly, in case of a set B of vertices, w.r.t. containment in all generator sets that support some element of B .
- (iii) w has *unique minimal support* if the supporting constituent $v\mathbb{C}\mathbb{G}[\alpha']$ in $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ with such minimal α' is unique (up to the obvious ambiguity in the choice of v within its α' -component in $\mathbb{H}[\alpha]$); in the case of a set B , the unique minimal supporting constituent $v\mathbb{C}\mathbb{G}[\alpha']$ is required to be contained in the minimal supporting constituent of any element of $B \cap \mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$.

Vertices of $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ have unique minimal supports by construction, viz. the intersection over all supporting constituents $v\mathbb{C}\mathbb{G}[\alpha']$, which is a supporting constituent (over the intersection of the corresponding generator sets). Much less obvious, but particularly important to us, is the case of β -components of small coset amalgams in the rôle of B in (ii). Unique minimal support for that case is part of the *cluster property* in Definition 7.6 below.

For trivial \mathbb{I} without edges, \mathbb{I} -skeletons consist of isolated vertices, and a small coset amalgam becomes a disjoint union of free amalgamation clusters. In that case, β -components B of free amalgamation clusters (when embedded into a 3-acyclic E-group) have unique minimal support: $\alpha_0 := \bigcap \{\alpha \subseteq E : \mathbb{G}[\alpha] \cap B \neq \emptyset\}$ is a minimal set of supporting generators, and $\mathbb{C}\mathbb{G}[\alpha_0] \subseteq \mathbb{C}\mathbb{G}$ a unique minimal support. We here try to lift this and related insights from free amalgamation clusters to small coset amalgams over non-trivial skeletons.

Definition 7.6. [cluster property]

Let $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ be defined according to Definition 7.1. This small coset amalgam has the *cluster property* if, for $\beta \subsetneq \alpha$, any β -component of $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ that does not directly intersect the embedded skeleton $\mathbb{H}[\alpha]$ has unique minimal

support, say $v\mathbb{C}\mathbb{G}[\alpha_0]$, isomorphic (as a β -graph) to a free amalgamation cluster of β' -cosets for $\beta' \subseteq \beta$ all containing a common $(\beta \cap \alpha_0)$ -coset of $v\mathbb{C}\mathbb{G}[\alpha_0]$.

The exception of the simpler case that the β -component in question intersects the skeleton, and hence coincides with a β -constituent of $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$, is necessary. While this β -constituent is uniquely determined, it cannot be *minimal* in the sense of Definition 7.5 since typical elements have smaller supports, some even empty support. In the non-degenerate case of a β -component that is disjoint from the skeleton, the cluster property states that the unique minimal supporting constituent (with a necessarily non-empty generator set α_0) serves as a unique bottleneck between the skeleton and that β -component. Moreover that β -component branches out, as a cluster, into β_i -cosets of constituents $v\mathbb{C}\mathbb{G}[\alpha_i] \supseteq v\mathbb{C}\mathbb{G}[\alpha_0]$, for all $\alpha_0 \subseteq \alpha_i \subsetneq \alpha$, from a single β_0 -coset of the minimal supporting constituent $v\mathbb{C}\mathbb{G}[\alpha_0]$ for $\beta_i := \beta \cap \alpha_i$.

There is an apparent similarity between the conditions expressed as bridge-freeness in Definition 7.3 (ii) and the cluster property as just defined. One may think of bridge-freeness as a condition that guarantees the analogue of the cluster property not just for β -components inside small coset amalgams $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$, but w.r.t. the intersections between $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ and full β -cosets in a surrounding $\mathbb{C}\mathbb{G}[\alpha]$.

Remark 7.7. *Let $\mathbb{G}[\alpha]$ be 3-acyclic. If the small coset amalgam $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ is defined according to Definition 7.1 and has the cluster property, then for $\beta, \gamma \subsetneq \alpha$, the β - and γ -components of elements w in $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ intersect as $\beta(w) \cap \gamma(w) = \delta(w)$ for $\delta := \beta \cap \gamma$; i.e. $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ is 2-acyclic.*

Proof. Let B and C be the vertex sets of the β - and γ -components of w in $\mathbb{K} := \mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$. The claim is trivial if both $\beta(w)$ and $\gamma(w)$ intersect $\mathbb{H}[\alpha]$. If just one of them, say B intersects $\mathbb{H}[\alpha]$, the other one has non-trivial minimal supporting constituent $v_0\mathbb{C}\mathbb{G}[\alpha_0]$ in $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$, and is given as a free amalgamation cluster of $\mathbb{G}[\gamma']$ -cosets for $\gamma' \subseteq \gamma$ centered on some $v_0g_0 \in v_0\mathbb{C}\mathbb{G}[\alpha_0]$. The cluster property implies that the minimal supporting constituent of $w \in C$ is of the form $v_0\mathbb{C}\mathbb{G}[\alpha'_0]$ for some α'_0 with $\alpha_0 \subseteq \alpha'_0 \subseteq \beta$. As $w \in v_0\mathbb{C}\mathbb{G}[\alpha'_0] \subseteq v_0\mathbb{C}\mathbb{G}[\beta]$, $B = v_0\mathbb{C}\mathbb{G}[\beta]$ follows. By 2-acyclicity of any one of the constituents of \mathbb{K} that contain $v_0g_0\mathbb{G}[\gamma']$ for any one of the γ' -cosets in the cluster C , the intersection of B with that γ' -coset is just $v_0g_0\mathbb{G}[\beta \cap \gamma']$,

In the remaining case, both B and C are disjoint from $\mathbb{H}[\alpha]$, hence correspond to free amalgamation clusters of $\mathbb{G}[\beta']$ - and $\mathbb{G}[\gamma']$ -cosets, respectively, centered on some $v_1g_1 \in v_1\mathbb{C}\mathbb{G}[\alpha_1]$ and $v_2g_2 \in v_2\mathbb{C}\mathbb{G}[\alpha_2]$ for minimal supporting constituents $v_1\mathbb{C}\mathbb{G}[\alpha_1]$ for B and $v_2\mathbb{C}\mathbb{G}[\alpha_2]$ for C . As w is contained in both, the cluster property now implies that its minimal supporting constituent is of the form $v_0\mathbb{C}\mathbb{G}[\alpha_0]$ such that $v_i\mathbb{C}\mathbb{G}[\alpha_i] \subseteq v_0\mathbb{C}\mathbb{G}[\alpha_0]$ for $i = 1, 2$. In particular $\alpha_1, \alpha_2 \subseteq \alpha_0$. Suppose now that these clusters meet not just in w but also in some w' so that $B = \beta(w) = \beta(w')$ and $C = \gamma(w) = \gamma(w')$; for the claim of the remark, it suffices to show that w and w' are δ -connected in \mathbb{K} . Just as for w , we find that the minimal supporting constituent for w' must be of the form $v'_0\mathbb{C}\mathbb{G}[\alpha'_0]$ such that $\alpha_1, \alpha_2 \subseteq \alpha'_0$. It follows that w and w' are elements of the constituents $v_0\mathbb{C}\mathbb{G}[\alpha_0]$ and $v'_0\mathbb{C}\mathbb{G}[\alpha'_0]$. These constituents $v_0\mathbb{C}\mathbb{G}[\alpha_0]$ and $v'_0\mathbb{C}\mathbb{G}[\alpha'_0]$

both contain the elements $v_1g_1 \in v_1\mathbb{C}\mathbb{G}[\alpha_1]$ and $v_2g_2 \in v_2\mathbb{C}\mathbb{G}[\alpha_2]$. W.r.t. to $\alpha_0 \cup \alpha'_0$, these two constituents form a free amalgam of copies of $\mathbb{C}\mathbb{G}[\alpha_0]$ and $\mathbb{C}\mathbb{G}[\alpha'_0]$, and w and w' are both β - and γ -connected inside this amalgam: β -connected via $v_1g_1 \in v_1\mathbb{C}\mathbb{G}[\alpha_1]$, and γ -connected via $v_2g_2 \in v_2\mathbb{C}\mathbb{G}[\alpha_2]$. As $\mathbb{G}[\alpha]$ is 3-acyclic, the free amalgam in question embeds isomorphically as an induced substructure into $\mathbb{C}\mathbb{G}[\alpha]$. By 2-acyclicity of $\mathbb{G}[\alpha]$, w and w' are δ -connected inside this free amalgam, hence also in \mathbb{K} . \square

The following can serve to lift the cluster property for small coset amalgams at level Γ_α , i.e. for $\alpha' \subsetneq \alpha$, to level α . It does, however, crucially require bridge-freeness at level Γ_α . That will necessitate a second lifting process, geared towards the freeness criteria that concern the embedding of $\mathbb{H}[\alpha]$ in $\mathbb{G}[\alpha]$, rather than intrinsic properties of $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$.

Recall from Observation 7.4 that under the assumptions of the lemma, which concern level Γ_α , the small coset amalgam $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ at level α is well-defined according to Definition 7.1.

Lemma 7.8. *Let $\mathbb{G}[\alpha]$ be 3-acyclic and such that, for all $\alpha' \subsetneq \alpha$, the embedded $\mathbb{I}[\alpha']$ -skeletons $\mathbb{H}[\alpha'] \subseteq_w \mathbb{C}\mathbb{G}[\alpha']$ are bridge-free and such that the embedded small coset amalgams $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha'], \mathbb{G}, \alpha')$ have the cluster property for all $\alpha' \subsetneq \alpha$. Then the small coset amalgam $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ has the cluster property.*

Proof. Let $V \subseteq \mathbb{G}[\alpha]$ be the vertex set of $\mathbb{K} = \mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$, $B \subseteq V$ the vertex set of the β -component of some $w \in V$ in \mathbb{K} . As $w \in \mathbb{K}$, some $\alpha' \subsetneq \alpha$ supports w , say $w = vg \in v\mathbb{C}\mathbb{G}[\alpha']$. For the case of interest we assume that B does not intersect the skeleton $\mathbb{H}[\alpha]$ of \mathbb{K} .

Suppose first that, for some constituent $v\mathbb{C}\mathbb{G}[\alpha']$ with $|\alpha'| = |\alpha| - 1$ (maximal in $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$) that supports B , the set $B' := B \cap v\mathbb{C}\mathbb{G}[\alpha'] \neq \emptyset$ is not supported by any strict subset of the generator set α' . We show that in this case $w = vg \in v\mathbb{C}\mathbb{G}[\alpha']$ implies $B = B' = vg\mathbb{C}\mathbb{G}[\beta']$ for $\beta' := \beta \cap \alpha'$. Otherwise some β -edge would have to link some vertex in B' to some vertex outside $v\mathbb{C}\mathbb{G}[\alpha']$; w.l.o.g. let these vertices be $w = vg$ and $v' \in B \setminus B'$, linked by a b -edge, $b \in \beta \setminus \alpha'$, stemming from some $\mathbb{C}\mathbb{G}[\alpha'']$ -constituent; but then $b \in \alpha'' \setminus \alpha'$ implies that $w \in B'$ would be supported by $\alpha' \cap \alpha'' \subsetneq \alpha'$, contradicting the assumption. So $B = B' = vg\mathbb{G}[\beta'] \subseteq v\mathbb{C}\mathbb{G}[\alpha']$. Uniqueness of $v\mathbb{C}\mathbb{G}[\alpha']$ as the minimal supporting constituent follows.

For the remaining cases, where B contains elements supported by smaller generator sets, all its non-empty intersections with constituents $v_i\mathbb{C}\mathbb{G}[\alpha_i]$ of \mathbb{K} for $|\alpha_i| = |\alpha| - 1$ will be β_i -cosets for $\beta_i := \beta \cap \alpha_i$ supported by generator sets $\alpha_{0i} \subsetneq \alpha_i$. Bridge-freeness for the embedded small coset amalgam $\mathbb{C}\mathbb{E}(\mathbb{H}[\alpha_i], \mathbb{G}, \alpha_i)$ within this $\mathbb{C}\mathbb{G}[\alpha_i]$ -constituent implies that the β_i -coset in question has a unique minimal supporting sub-constituent $v_i\mathbb{C}\mathbb{G}[\alpha_{0i}]$ for $\alpha_{0i} \subsetneq \alpha_i$ (after shifting v_i within its α_i -component of $\mathbb{H}[\alpha]$ if necessary, to avoid unnecessary distinctions). The β -component B therefore is the union of β_i -cosets of the form

$$v_i g_i \mathbb{G}[\beta_i] \subseteq v_i \mathbb{C}\mathbb{G}[\alpha_i] \text{ for } g_i \in \mathbb{G}[\alpha_{0i}], \alpha_{0i} \subsetneq \alpha_i \subsetneq \alpha, |\alpha_i| = |\alpha| - 1.$$

We need to show that these parts overlap in the form of a free amalgamation cluster centered on some β' -coset of a unique minimal supporting constituent for B in \mathbb{K} . Consider two such β_i -cosets $v_i g_i \mathbb{G}[\beta_i] \subseteq v_i \mathbb{C}\mathbb{G}[\alpha_i]$ with $g_i \in \mathbb{G}[\alpha_{0i}]$, say for $i = 1, 2$, that overlap in \mathbb{K} . A common vertex $w \in v_1 g_1 \mathbb{G}[\beta_1] \cap v_2 g_2 \mathbb{G}[\beta_2]$ is of the form $v_0 g_0$ where $g_0 \in \mathbb{G}[\alpha_0]$ for a minimal supporting set α_0 of generators for w , so that in particular $\alpha_0 \subseteq \alpha_1 \cap \alpha_2$. (The situation is as represented in Figure 4 on the right.) Looking at $v_i g_i$ and $v_0 g_0$, which are β_i -connected in $v_i \mathbb{C}\mathbb{G}[\alpha_i]$, we can invoke bridge-freeness in this $\mathbb{C}\mathbb{G}[\alpha_i]$ -constituent to find $v_0 i g_0 i \in v_i g_i \mathbb{G}[\beta_i]$ where $g_0 i \in \mathbb{G}[\alpha'_{0i}]$ for $\alpha'_{0i} \subseteq \alpha_{0i} \cap \alpha_0$, by minimality of those two generator sets.⁹ So the cluster property for this $\mathbb{C}\mathbb{G}[\alpha_i]$ -constituent implies that $v_i \mathbb{C}\mathbb{G}[\alpha_{0i}] \subseteq \alpha_0$. The cluster property for the $\mathbb{C}\mathbb{G}[\alpha_0]$ -constituent implies that $v_1 \mathbb{C}\mathbb{G}[\alpha_{01}] = v_2 \mathbb{C}\mathbb{G}[\alpha_{02}] \subseteq v_0 \mathbb{C}\mathbb{G}[\alpha_0]$ and that $v_1 g_1 \mathbb{G}[\beta_1]$ and $v_2 g_2 \mathbb{G}[\beta_2]$ overlap in $v_0 \mathbb{C}\mathbb{G}[\alpha_0]$. So any two directly overlapping β_i -parts of B share the same unique minimal supporting constituent in their $\mathbb{C}\mathbb{G}[\alpha_i]$ -constituents (for maximal $\alpha_i \subsetneq \alpha$) and overlap inside that minimal constituent. This shared minimal constituent $v_0 \mathbb{C}\mathbb{G}[\alpha_0]$ is therefore minimal and unique also in \mathbb{K} . The nature of B as a union of β_i -cosets that pairwise overlap in this minimal constituent implies that it forms a free amalgamation cluster. \square

It follows that, under suitable conditions, the classes of free amalgamation clusters, small coset amalgams, and of free amalgams between them, satisfy useful closure conditions w.r.t. passage to β -components for smaller β . For the second part of the following corollary we use 2-acyclicity of $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ from Remark 7.7 and appeal to Observation 4.4. We also note that any β -component that intersects the skeleton of any small coset amalgam with constituents from $\Gamma_k(\mathbb{G})$ and for $\beta \in \Gamma_k$ is necessarily isomorphic to a constituent $\mathbb{C}\mathbb{G}[\beta]$ -copy, and cannot possibly extend beyond that in any E-graph extensions.

Corollary 7.9. *If $\mathbb{G}[\alpha]$ is 3-acyclic and for all $\alpha' \subsetneq \alpha$ the $\mathbb{H}[\alpha']$ -skeletons $\mathbb{H}[\alpha'] \subseteq_w \mathbb{C}\mathbb{G}[\alpha']$ are bridge-free and their embedded small coset amalgams have the cluster property, then the small coset amalgam $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ is well-defined, has the cluster property and is 2-acyclic. Moreover, all the free amalgams*

$$\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha), v \oplus \mathbb{C}\mathbb{G}[\alpha']$$

are well-defined for all vertices v of $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ and all $\alpha' \subsetneq \alpha$. Their β -components, for $\beta \subsetneq \alpha$, are isomorphic to $\mathbb{C}\mathbb{G}[\beta]$ itself or to a free amalgam between a free amalgamation cluster over $\Gamma(\beta)$ and some $\mathbb{C}\mathbb{G}[\beta']$ for $\beta' \subsetneq \beta$ (including, for $\beta' = \emptyset$, plain free amalgamation clusters over $\Gamma(\beta)$).

We note that this corollary establishes downward compatibility in the sense that the passage to small coset amalgams and their free amalgams with individual small $\mathbb{C}\mathbb{G}[\alpha']$ has the desired downward closure property: the functor that associates free amalgamation clusters and their free amalgams with individual $\mathbb{C}\mathbb{G}[\alpha']$ is conservative in the sense of Definition 2.13 (as established

⁹We invoke bridge-freeness rather than the cluster property here, since it is not clear that $v_i g_i$ and $v_0 g_0$ are β_i -connected in $\text{CE}(\mathbb{H}[\alpha_i], \mathbb{G}, \alpha_i)$, even though they are β_i -connected in $v_i \mathbb{C}\mathbb{G}[\alpha_i]$.

in Corollaries 4.14 and 4.16). By the above corollary the isomorphism types of the β -components of extended small coset amalgams are determined by $\Gamma_\alpha(\mathbb{C}\mathbb{G})$, across all skeletons $\mathbb{H}[\alpha]$ that support the construction of these small coset amalgams according to Definition 7.1. The underlying conservative functor \mathbf{F} that produces all these isomorphism types in the context of Corollaries 4.14 and 4.16 therefore also safeguards the inclusion of extended small coset amalgams as components in \mathbf{E} -graphs \mathbb{K} for inductive unfolding steps $\mathbb{G} \preceq_k \hat{\mathbb{G}} = \text{sym}(\mathbb{K})$ when dealing with $|\alpha| = k$, i.e. $\alpha \in \Gamma_{k+1}$.

It requires one more argument to lift the (intrinsic) cluster property for $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ to bridge-freeness (in relation to some surrounding \mathbf{E} -group $\hat{\mathbb{G}}$ that unfolds \mathbb{G} (and hence also embedded $\mathbb{I}[\alpha]$ -skeletons) while preserving all structure at the level of Γ_α . This is the purpose of the following lemma.

Lemma 7.10. *Suppose that $\mathbb{G}[\alpha]$ is 3-acyclic and such that all embedded $\mathbb{I}[\alpha']$ -skeletons $\mathbb{H}[\alpha'] \subseteq_w \mathbb{C}\mathbb{G}[\alpha']$ for $\alpha' \subsetneq \alpha$ are bridge-free. Let $\hat{\mathbb{G}} \succ_\alpha \mathbb{G}$ be 3-acyclic and compatible with all α -graphs of the form*

$$\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha), v \oplus \mathbb{C}\mathbb{G}[\alpha'] \text{ for } \alpha' \subsetneq \alpha.$$

Then any embedded $\mathbb{I}[\alpha]$ -skeleton $\hat{\mathbb{H}}[\alpha] \subseteq_w \mathbb{C}\hat{\mathbb{G}}[\alpha]$ is bridge-free.

Note that $\text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha) \simeq \text{CE}(\hat{\mathbb{H}}[\alpha], \mathbb{G}, \alpha)$ as $\hat{\mathbb{G}} \succ_\alpha \mathbb{G}$, but $\hat{\mathbb{H}}[\alpha]$ may be a non-trivial unfolding of $\mathbb{H}[\alpha]$. By the previous lemma, $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ has the cluster property, and the amalgams $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha), v \oplus \mathbb{C}\mathbb{G}[\alpha']$ above are well-defined. The conclusion of the lemma holds for any $\hat{\mathbb{G}}' \succ_k \hat{\mathbb{G}}$ if $\alpha \in \Gamma_k$.

By Observation 7.4 the conclusion of the lemma implies that the small coset amalgam $\text{CE}(\hat{\mathbb{H}}, \hat{\mathbb{G}}, \alpha)$ embeds as an induced substructure into $\mathbb{C}\hat{\mathbb{G}}$ and, by Lemma 7.8, it also has the cluster property.

Proof. We first observe that $\mathbb{H}[\alpha'] \simeq \hat{\mathbb{H}}[\alpha']$ for $\alpha' \subsetneq \alpha$, due to $\hat{\mathbb{G}} \succ_\alpha \mathbb{G}$. Together with 2-acyclicity of $\hat{\mathbb{G}}$, this further implies that the embedded $\mathbb{I}[\alpha]$ -skeleton $\hat{\mathbb{H}}[\alpha]$ is 2-acyclic. Any α_i -components $\alpha_i(v)$ for $\alpha_1, \alpha_2 \subsetneq \alpha$ in $\hat{\mathbb{H}}[\alpha]$ that also meet in a vertex v' of $\hat{\mathbb{H}}[\alpha]$ outside $\alpha_0(v)$, for $\alpha_0 := \alpha_1 \cap \alpha_2$, would establish a violation of freeness for $\mathbb{H}[\alpha_1] \simeq \hat{\mathbb{H}}[\alpha_1]$ w.r.t. $\alpha_0 \subsetneq \alpha_1$ in $\mathbb{C}\mathbb{G}[\alpha_1]$.

Claim 1. $\hat{\mathbb{H}}[\alpha] \subseteq_w \mathbb{C}\hat{\mathbb{G}}[\alpha]$ is free in $\mathbb{C}\hat{\mathbb{G}}$, so that $\text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha)$ embeds isomorphically as a weak substructure of $\mathbb{C}\hat{\mathbb{G}}[\alpha] \subseteq \mathbb{C}\hat{\mathbb{G}}$.

This relies on compatibility of $\hat{\mathbb{G}}$ with the \mathbf{E} -graphs of the form $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ (which are included in the stated collection for $\alpha' = \emptyset$). For the following compare Figure 5. Suppose that for $\alpha_i \subsetneq \alpha$ and v_1, v_2 from $\hat{\mathbb{H}}[\alpha]$ there is some $g \in v_1\mathbb{G}[\alpha_1] \cap v_2\mathbb{G}[\alpha_2]$ in $\hat{\mathbb{G}}$. Let $g_i \in \mathbb{G}[\alpha_i]$ and $u \in \alpha^*[\mathbb{I}]$ be such that $v_2 = v_1[u]_{\hat{\mathbb{G}}}$ and $g = v_i g_i$ for $i = 1, 2$. It follows that

$$g_1 = [u]_{\hat{\mathbb{G}}} g_2 \text{ in } \hat{\mathbb{G}},$$

which by compatibility transfers to the operation of $\hat{\mathbb{G}}$ on $\mathbb{K} := \text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$. So there are vertices v'_1, v'_2 in $\mathbb{H}[\alpha] \subseteq \mathbb{C}\mathbb{G}[\alpha]$ such that $v'_2 = v'_1[u]_{\mathbb{K}}$ with a




Figure 5: Walks in relevant α_i -components of $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ and $\text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha)$.

non-trivial intersection of the constituents $v'_i \mathbb{C}\mathbb{G}[\alpha_i]$ in $v' = v'_1 g_1 = v'_2 g_2$. By definition of $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$, the α_i -components of the v'_i overlap in the skeleton $\mathbb{H}[\alpha]$, say in v'_0 . Let $v'_1 = v'_0 [u_{01}]_{\mathbb{K}}$ and $v'_2 = v'_0 [u_{02}]_{\mathbb{K}}$ for $u_{0i} \in \alpha_i^*[\mathbb{I}]$, and $v' = v'_0 g_0$ for some $g_0 \in \mathbb{G}[\alpha_0]$ where $\alpha_0 = \alpha_1 \cap \alpha_2$. Then

$$[u_{01}]_{\mathbb{G}} g_1 = g_0 = [u_{02}]_{\mathbb{G}} g_2$$

in \mathbb{G} . This transfers to $\hat{\mathbb{G}}$, since both \mathbb{G} and $\hat{\mathbb{G}}$ are 2-acyclic, and $\mathbb{C}\mathbb{G}[\alpha_i] \simeq \mathbb{C}\hat{\mathbb{G}}[\alpha_i]$, which implies that free amalgams between such constituents produce isomorphic results in $\mathbb{C}\mathbb{G}$ and $\mathbb{C}\hat{\mathbb{G}}$. So also in $\hat{\mathbb{G}}$

$$g_1 g_2^{-1} = [u_{01}]_{\hat{\mathbb{G}}}^{-1} [u_{02}]_{\hat{\mathbb{G}}}$$

is represented by a compositions of an $\mathbb{I}[\alpha_1]$ -walk with an $\mathbb{I}[\alpha_2]$ -walk inside the embedded skeleton $\hat{\mathbb{H}}[\alpha]$. In particular, as $g_1 g_2^{-1}$ connects v_1 to v_2 in $\text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha)$, $v_0 := v_1 [u_{01}]_{\hat{\mathbb{G}}}^{-1}$ witnesses the non-empty overlap between $\alpha_1(v_1)$ and $\alpha_2(v_2)$ in the skeleton $\hat{\mathbb{H}}[\alpha]$. Moreover, $g = v_1 [u_{01}]_{\hat{\mathbb{G}}}^{-1} g_0$ is α_0 -connected to this overlap by an α_0 -walk. This implies freeness of $\hat{\mathbb{H}}[\alpha]$ in $\mathbb{C}\hat{\mathbb{G}}$ and that $\text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha)$ embeds into $\mathbb{C}\hat{\mathbb{G}}$, so far as a weak substructure.

Claim 2. $\hat{\mathbb{H}}[\alpha] \subseteq_w \mathbb{C}\hat{\mathbb{G}}$ is bridge-free.¹⁰

For Claim 2 we need to show that elements of $\text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha) \subseteq_w \mathbb{C}\hat{\mathbb{G}}$ that are members of the same β -coset of $\hat{\mathbb{G}}$ for some $\beta \subsetneq \alpha$ must be β -connected inside $\text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha)$ in the manner expressed in criterion (*) of Definition 7.3, as illustrated in Figure 4. In the following we use $\mathbb{K} := \text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ to denote the small coset amalgam at the level of \mathbb{G} , and $\hat{\mathbb{K}} := \text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha) \subseteq_w \mathbb{C}\hat{\mathbb{G}}$ for the small coset amalgam at the level of $\hat{\mathbb{G}}$.

¹⁰This also strengthens Claim 1 to $\text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha) \subseteq \mathbb{C}\hat{\mathbb{G}}$, cf. Observation 7.4.




Figure 6: The argument for bridge-freeness.

Consider elements $v_i g_i \in v_i \mathbb{C}\hat{\mathbb{G}}[\alpha_i]$ in $\hat{\mathbb{K}}$ for v_i from $\hat{\mathbb{H}}[\alpha] \subseteq_w \mathbb{C}\hat{\mathbb{G}}[\alpha]$, represented as members of constituents $v_i \mathbb{C}\hat{\mathbb{G}}[\alpha_i]$ over minimal supporting sets $\alpha_i \subsetneq \alpha$ (existence and uniqueness of minimal supports for elements is intrinsically guaranteed by definition of $\text{CE}(\hat{\mathbb{H}}[\alpha], \hat{\mathbb{G}}, \alpha)$; here we could meanwhile also argue by freeness of $\hat{\mathbb{H}}[\alpha]$).

Assume that $v_2 g_2$ is linked to $v_1 g_1$ by some β -walk, i.e. that $v_2 g_2 = v_1 g_1 h$ for some $h \in \hat{\mathbb{G}}[\beta]$, say $h = [w]_{\hat{\mathbb{G}}}$ for $w \in \beta^*$. Let $v_2 = v_1 [u]_{\hat{\mathbb{G}}}$ for some $u \in \alpha^*[\mathbb{I}]$ tracing an $\mathbb{I}[\alpha]$ -walk from v_1 to v_2 in $\hat{\mathbb{H}}[\alpha]$. It follows that

$$g_1 [w]_{\hat{\mathbb{G}}} = [u]_{\hat{\mathbb{G}}} g_2 \text{ in } \hat{\mathbb{G}}.$$

Consider vertices v'_i in $\mathbb{H}[\alpha] \subseteq_w \mathbb{K} = \text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ that are linked by a u -labelled walk, such that in particular $v'_2 = v'_1 [u]_{\mathbb{K}}$ (cf. the lefthand side of Figure 6). By compatibility of $\hat{\mathbb{G}}$ with \mathbb{K} , the operation of $[w]_{\mathbb{K}}$ maps $v'_1 g_1$ to $v'_2 g_2$ in \mathbb{K} . Therefore $v'_1 g_1$ and $v'_2 g_2$ are β -connected in \mathbb{K} . As \mathbb{K} has the cluster property (cf. Lemma 7.8), the β -component of $v'_1 g_1$ and $v'_2 g_2$ in \mathbb{K} either intersects the skeleton $\mathbb{H}[\alpha] \subseteq_w \mathbb{K}$ or else has a unique minimal supporting constituent $v'_0 \mathbb{C}\mathbb{G}[\alpha_0]$. In the latter case, the cluster property yields some element $v'_0 g_0 \in v'_0 \mathbb{C}\mathbb{G}[\alpha_0]$, in which the cluster can be rooted; and as elements of this cluster the $v'_i g_i$ must be of the form $v'_i g_i \in v'_0 g_0 \mathbb{G}[\beta'_i]$ where $\beta_i := \beta \cap \alpha'_i$ for suitable $\alpha'_i \subsetneq \alpha$ with $\alpha_i \subseteq \alpha'_i$ (generator sets α_i are minimal for $v_i g_i$ in $\hat{\mathbb{K}}$, and hence also for $v'_i g_i$ in \mathbb{K}). In the former case, i.e. if the β -cluster degenerates to a β -constituent of the form $v'_0 \mathbb{C}\mathbb{G}[\beta]$, we may use $v'_0 g_0 = v'_0 1 \in \mathbb{C}\mathbb{G}[\alpha_0]$ for $\alpha_0 = \emptyset$, and similarly obtain that $v'_i g_i \in v'_0 g_0 \mathbb{G}[\beta'_i]$ where $\beta_i := \beta \cap \alpha'_i$ for suitable $\alpha'_i \subsetneq \alpha$ with $\alpha_i \subseteq \alpha'_i$ (in this case, indeed, minimality of α_i implies that $\alpha_i \subseteq \beta$ so that $\beta_i = \alpha_i = \alpha'_i$ can be used). In either case, therefore, $v'_i g_i = v'_0 g_0 [w_i]_{\mathbb{G}}$ in \mathbb{K} , for suitable $w_i \in \beta_i^*$. And as the operation of $[w]_{\mathbb{K}}$ maps $v'_1 g_1$ to $v'_2 g_2$ in \mathbb{K} , $v'_0 g_0$ is a fixpoint of the operation of

$$[w_1]_{\mathbb{K}} [w]_{\mathbb{K}} [w_2^{-1}]_{\mathbb{K}}.$$

The same argument applies to $v'_0 g_0$ as an element of the free amalgam of \mathbb{K} with a copy of $\mathbb{C}\mathbb{G}[\beta]$ attached at $v'_0 g_0$,

$$\mathbb{K}, v'_0 g_0 \oplus \mathbb{C}\mathbb{G}[\beta].$$

Because $\hat{\mathbb{G}}$ is assumed to be also compatible with that, also here the operation of $[w]$ must map $v'_1 g_1$ to $v'_2 g_2$. In $\mathbb{C}\mathbb{G}[\beta] \subseteq \mathbb{K}, v'_0 g_0 \oplus \mathbb{C}\mathbb{G}[\beta]$, however, $w_1 w w_2^{-1} \in \beta^*$ can only have a fixpoint if $[w_1 w w_2^{-1}]_{\mathbb{G}} = 1$. So $[w]_{\mathbb{G}} = [w_1]_{\mathbb{G}}^{-1} [w_2]_{\mathbb{G}}$ and therefore also $[w]_{\hat{\mathbb{G}}} = [w_1]_{\hat{\mathbb{G}}}^{-1} [w_2]_{\hat{\mathbb{G}}}$ in $\hat{\mathbb{G}}$, since $\mathbb{G} \preceq_{\alpha} \hat{\mathbb{G}}$ and $\beta \not\subseteq \alpha$.

It follows that also in $\hat{\mathbb{G}}$ there is $v_0 \in \hat{\mathbb{H}}[\alpha] \subseteq_w \hat{\mathbb{K}}$ such that

$$v_1 g_1 [w_1]_{\hat{\mathbb{G}}}^{-1} = v_0 g_0 = v_2 g_2 [w_2]_{\hat{\mathbb{G}}}^{-1}$$

so that $v_1 g_1$ and $v_2 g_2$ are β -related to $v_0 g$ within $v_1 \mathbb{C}\mathbb{G}[\alpha'_1]$ and $v_2 \mathbb{C}\mathbb{G}[\alpha'_2]$, respectively, as required by bridge-freeness. \square

8 Construction of N-acyclic E-groups over I

Again, we fix a constraint graph \mathbb{I} and only consider E-groups \mathbb{G} that are compatible with \mathbb{I} according to Proviso 6.1. We collect the desirable qualities of E-groups towards freeness and cluster properties at levels defined in terms of the sets Γ_k , i.e. w.r.t. the numbers of generators from E that are involved. This is geared towards an inductive construction for successively increasing the level k , see Lemma 8.2 and Figure 7. In the inductive process we may assume global N -acyclicity (for the target $N \geq 3$) for the essential intermediate E-groups \mathbb{G} , by interleaving uses of Proposition 5.2.

Definition 8.1. Let $N \geq 3$. For $k \geq 1$, \mathbb{G} is N -good for \mathbb{I} at level Γ_k if

- (i) \mathbb{G} is N -acyclic,
- (ii) the embedded \mathbb{I} -skeletons $\mathbb{H}[\alpha]$ are bridge-free in $\mathbb{C}\mathbb{G}[\alpha]$, for $\alpha \in \Gamma_k$,
- (iii) \mathbb{G} is compatible with
 - amalgamation chains of lengths up to $N - 2$ over $\Gamma_k(\mathbb{C}\mathbb{G})$,
 - free amalgamation clusters over $\Gamma_k(\mathbb{C}\mathbb{G})$, and with
 - free amalgams of such clusters with an extra $\mathbb{C}\mathbb{G}[\alpha'] \in \Gamma_k(\mathbb{C}\mathbb{G})$.

N -goodness at level Γ_1 is just N -acyclicity: conditions (ii) and (iii) are vacuous for $\alpha = \emptyset$. Goodness at level $\Gamma_{|E|+1}$ is the eventual goal. The following lemma allows us to increase the level of N -goodness by one, so that inductively we obtain an N -acyclic \mathbb{G} that is (bridge-)free over its embedded $\mathbb{I}[\alpha]$ -skeletons for all $\alpha \subseteq E$; compare Figure 7.

Lemma 8.2. Let $N \geq 3$, and let \mathbb{G} be N -good for \mathbb{I} at level Γ_k . Let $\hat{\mathbb{G}}^*$ be obtained as $\hat{\mathbb{G}}^* := \text{sym}(\mathbb{K})$ for \mathbb{K} with components $\mathbb{C}\mathbb{G}$ and, for all embedded \mathbb{I} -skeletons $\mathbb{H}[\alpha] \subseteq_w \mathbb{C}\mathbb{G}[\alpha]$ where $|\alpha| = k$, i.e. for $\alpha \in \Gamma_{k+1} \setminus \Gamma_k$:

$$\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha), v \oplus \mathbb{C}\mathbb{G}[\alpha'] \text{ for all } \alpha' \not\subseteq \alpha.$$

Then $\hat{\mathbb{G}}^* \succ_k \mathbb{G}$ and any $\hat{\mathbb{G}} \succ_{k+1} \hat{\mathbb{G}}^*$ that is N -acyclic and compatible with free amalgamation clusters over $\Gamma_{k+1}(\hat{\mathbb{G}}^*)$ and free amalgams of such clusters with an extra $\mathbb{C}\mathbb{G}[\alpha'] \in \Gamma_{k+1}(\mathbb{C}\hat{\mathbb{G}}^*)$ is N -good for \mathbb{I} at level Γ_{k+1} .

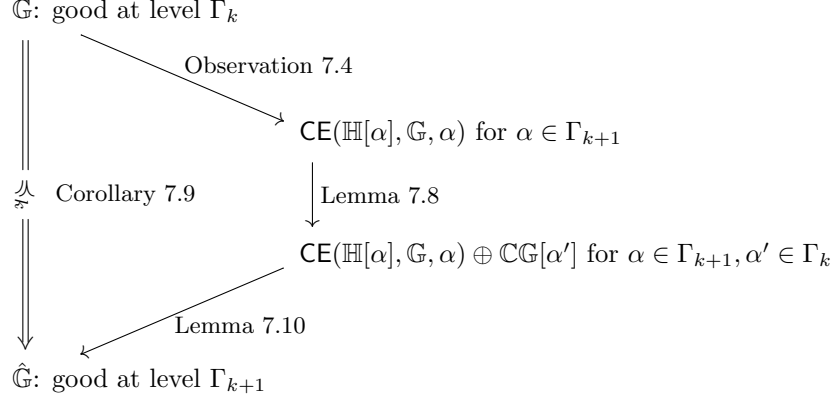


Figure 7: Inductive pattern: from level Γ_k to level Γ_{k+1} for Lemma 8.2.

Proof. Small coset amalgams $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha)$ for $\alpha \in \Gamma_{k+1}$ as well as their free amalgams $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha), v \oplus \mathbb{CG}[\alpha']$ with $\mathbb{CG}[\alpha']$ for $\alpha' \subsetneq \alpha$ are well-defined by Corollary 7.9. These ingredients in \mathbb{K} for $|\alpha| = k$ together with \mathbb{G} being good for \mathbb{I} at level Γ_k , guarantees that \mathbb{G} itself is compatible with all β -components of \mathbb{K} for $\beta \in \Gamma_k$, by Corollary 7.9. Therefore \mathbb{CG} as an ingredient in \mathbb{K} entails that $\hat{\mathbb{G}}^* = \text{sym}(\mathbb{K})$ unfolds \mathbb{G} in a manner that preserves $\Gamma_k(\mathbb{G})$, i.e. such that $\mathbb{G} \preceq_k \hat{\mathbb{G}}^*$. It essentially remains to argue that the embedded α -skeletons $\hat{\mathbb{H}}^*[\alpha] \subseteq_w \mathbb{CG}^*[\alpha]$ for $|\alpha| = k$ are bridge-free: condition (ii) for goodness at level Γ_{k+1} boils down to the case for $|\alpha| = k$ as this condition is preserved by \preceq_k for smaller α . But compatibility of $\hat{\mathbb{G}}^*$ with the $\text{CE}(\mathbb{H}[\alpha], \mathbb{G}, \alpha), v \oplus \mathbb{CG}[\alpha']$, which are part of \mathbb{K} for these α and $\alpha' \subsetneq \alpha$, guarantees bridge-freeness by Lemma 7.10.

It is clear that $\hat{\mathbb{G}} \succ_{k+1} \hat{\mathbb{G}}^*$ is good at level Γ_{k+1} whenever N -acyclicity for condition (i) and compatibility conditions (iii) for goodness at level Γ_{k+1} are met. This can be achieved in a \preceq_{k+1} -preserving unfolding chain, since already $\mathbb{G} \preceq_k \hat{\mathbb{G}}^*$ satisfies these conditions at level Γ_k , so that the new compatibility requirements can be implemented without affecting subgroups over Γ_{k+1} . This uses the conservative nature of the functors (cf. Definition 2.13) that generate amalgamation chains, free amalgamation clusters and amalgams of free amalgamation clusters with one extra $\mathbb{CG}[\alpha]$:

- Corollary 4.9 for amalgamation chains,
- Corollaries 4.14 and 4.16 for plain and extended amalgamation clusters.

□

Similar to Theorem 5.3 we sum up the construction and emphasise its preservation of symmetries over E . For N -acyclicity over \mathbb{I} compare Definition 6.9 and its connection with freeness from Lemma 6.17.

Theorem 8.3. *For every finite E-group \mathbb{G} that is compatible with the constraint graph \mathbb{I} and every $N \geq 2$ there is a finite E-group $\hat{\mathbb{G}} \succ \mathbb{G}$ that is N -acyclic over \mathbb{I} and fully symmetric over \mathbb{G} in the sense that every permutation $\rho \in \text{Sym}(E)$ 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 $\mathbb{G} := \text{sym}(\mathbb{H})$ where \mathbb{H} is the disjoint union of \mathbb{I} and the hypercube $\mathbb{H} = 2^E$.

9 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 $\pi_e \in \text{Sym}(V)$ 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 § 6, if we think of S -markings, e.g. in \mathbb{I} -skeletons according to Definition 6.10, as marking out distinct sorts of vertices. 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 are very naturally cast in terms of multi-sorted multi-graph structures and related groupoids, cf. [14, 16]. In this section we directly reduce the construction of groupoids with the desired coset acyclicity properties to the construction of groups from the previous sections.

9.1 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 or monoids 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 9.1. [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 $\iota_i: E \rightarrow S$ 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 $E[s, s'] := \{e \in E : \iota_1(e) = s, \iota_2(e) = s'\}$ 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. 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 9.2. [\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 $s = s_0, e_1, s_1, \dots, e_n, s_n = t$ such that $\iota_1(e_i) = s_{i-1}$ and $\iota_2(e_i) = s_i$ for $1 \leq i \leq n$, this walk is fully determined by the sequence of edges and can be identified with the word $w = e_1 \dots e_n \in E^*$. So we think of $E^*[\mathbb{I}]$ as the set of all words $w = e_1 \dots e_n$ with $\iota_2(e_i) = \iota_1(e_{i+1})$ for $1 \leq i < n$, 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

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

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 *sorts* marked by *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.

9.2 I-groupoids and their Cayley graphs

Definition 9.3. [\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;¹¹
- (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 $h: \hat{\mathbb{G}} \rightarrow \mathbb{G}$ (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 9.4. *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, 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 2.15 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 automorphism 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 9.5. [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 \mathbb{E} -groupoid, $\mathbb{G}^\rho \simeq \mathbb{G}$.

¹¹Some of the sets $G_{s,t}$ may be empty as \mathbb{I} is not required to be connected.

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

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

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 the \mathbb{I} -groupoid \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 9.6. [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 $\mathbb{C}\mathbb{G} := \text{Cayley}(\mathbb{G}) = (G, (R_e)_{e \in E})$ with vertex set G and edge sets of colour $e \in E$ according to

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

As with Cayley graphs for \mathbf{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 as individual constants 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 -edges 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 [12] 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 9.7. *The \mathbb{I} -groupoid \mathbb{G} is isomorphic to the \mathbb{I} -groupoid generated by the following bijections π_e for $e \in E[s, s']$:*

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

where $g \cdot g_e$ is identified as the unique vertex g' of $\mathbb{C}\mathbb{G}$ for which $(g, g') \in R_e$. Here $1_s = \text{id}_{G[* , s]}$ 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.¹²

Definition 9.8. [\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 $\pi_e: V_{\iota_1(e)} \rightarrow V_{\iota_2(e)}$). 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, which also shares every symmetry of the \mathbb{I} -groupoid (cf. Definition 9.5 for those symmetries). Conversely, any complete \mathbb{I} -graph determines an \mathbb{I} -groupoid in the manner indicated for this special case in Observation 9.7 above. In a complete \mathbb{I} -graph \mathbb{H} as in Definition 9.8, the composition of the π_e along $w = e_1 \cdots e_n \in E^*[\mathbb{I}, s, t]$, $\pi_w = \prod_{i=1}^n \pi_{e_i} = \pi_{e_n} \circ \cdots \circ \pi_{e_1}$, induces a bijection $\pi_w: V_s \rightarrow V_t$, 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 $G = \{\pi_w: w \in E^*[\mathbb{I}]\}$.

Definition 9.9. [$\text{sym}(\mathbb{H})$]

From a complete \mathbb{I} -graph \mathbb{H} as in Definition 9.8, with induced partial bijections π_w for $w \in E^*[\mathbb{I}]$, we obtain the \mathbb{I} -groupoid

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

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 9.7 can be restated as $\text{sym}(\mathbb{C}\mathbb{G}) \simeq \mathbb{G}$ if \mathbb{G} is an \mathbb{I} -groupoid with Cayley graph $\mathbb{C}\mathbb{G}$.

Definition 9.10. [compatibility]

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

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

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

¹²Indeed the proposal of a naive completion operation accounts for the major flaw in [15].

Observation 9.11. *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 $\mathbb{C}\mathbb{G}$ of \mathbb{G} if, and only if, $\hat{\mathbb{G}} \succcurlyeq \mathbb{G}$, if, and only if the map $h: \hat{\mathbb{G}} \rightarrow \mathbb{G}$ which maps $[w]_{\hat{\mathbb{G}}}$ to $[w]_{\mathbb{G}}$ is well-defined (and thus the homomorphism in question).*

9.3 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 9.12. [coset cycles]

Let \mathbb{G} be an \mathbb{I} -groupoid, $n \geq 2$. 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 9.13. [N -acyclicity]

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

10 Construction of N -acyclic \mathbb{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 $\hat{\mathbb{E}}$ -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 $\mathbb{C}\mathbb{G}$ within the direct product of $\hat{\mathbb{I}}$ with the Cayley graph $\mathbb{C}\hat{\mathbb{G}}$ of an $\hat{\mathbb{E}}$ -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 $e \in E$ three 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


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

we represent directed e -(multi-)edges as walks of length 3 in an $\hat{\mathbb{E}}$ -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 $s_{e'}$:

$$s \xrightarrow{\{e\}} s_e \xrightarrow{\{e, e'\}} s_{e'} \xrightarrow{\{e'\}} s'$$

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{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{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{E} -groups and their Cayley graphs. Overall, this will allow us to directly extract \mathbb{I} -groupoids from suitable \hat{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

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

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

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

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 $\hat{\alpha}$ 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:

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

10.1 Groupoids from groups

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

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

is an \hat{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{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 $\mathbb{G} := \mathbb{G}[\mathbb{I}]$ 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{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 $\{(s, [\hat{w}]_{\hat{\mathbb{G}}}) \in \hat{\mathbb{I}} \otimes \hat{\mathbb{G}} : \hat{w} \in \hat{E}^*[\hat{\mathbb{I}}, s, t]\}$ for $s, t \in S \subseteq \hat{S}$, and put

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

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

$$G := \bigcup_{s,t \in S} G_{s,t}$$

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:

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

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

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

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 9.4) 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 $\hat{\mathbb{E}}$ -graph $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ that embed \mathbb{G} :

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

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 $\hat{\mathbb{E}}$ -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} :

$$\begin{aligned} \mathbb{H}_0 &= (G, (G_{*,s})_{s \in S}, (R_e)_{e \in E}) \quad \text{where, for } e = (s, s'), \\ R_e &= \{((t, [u]_{\hat{\mathbb{G}}}), (t, [u\{e\}\{e, e^{-1}\}\{e^{-1}\}]_{\hat{\mathbb{G}}})) : u \in \hat{E}^*[\hat{\mathbb{I}}, t, s]\}. \end{aligned}$$

Observation 10.1. *Let $\hat{\mathbb{G}}$ be an $\hat{\mathbb{E}}$ -group that is compatible with $\hat{\mathbb{I}}$. Then the Cayley graph $\mathbb{C}\hat{\mathbb{G}}$ 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 $\hat{\mathbb{H}}_0 \subseteq \hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$.*

10.2 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 10.2. *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{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{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 $\hat{\rho} \in \text{Sym}(\hat{E})$ 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 10.3. *In the situation of Proposition 10.2 there is a natural translation of coset cycles in the \mathbb{I} -groupoid $\mathbb{G} = \hat{\mathbb{G}}/\mathbb{I}$ based on the map $\hat{\cdot}$ 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

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

be a coset cycle in the groupoid \mathbb{G} , according to Definition 9.12, 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 $\hat{\cdot}$ -translation of an α_i -word of generators representing $g_i^{-1}g_{i+1} \in \mathbb{G}$. The natural $\hat{\cdot}$ -translation of the cycle $(*)$ into $\hat{\mathbb{G}}$ is

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

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 9.12 for $(*)$.

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

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

where $\alpha_{i, i\pm 1} = \alpha_i \cap \alpha_{i\pm 1}$.

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 $e \in \alpha_{i,i+1}$. \square

Theorem 10.4. *For every finite constraint pattern $\mathbb{I} = (S, E)$, every complete \mathbb{I} -graph \mathbb{H} and every $N \geq 2$ 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 $\mathbb{H} := \text{sym}(\mathbb{I})$ (regarding \mathbb{I} as a complete \mathbb{I} -graph according to Definition 9.8) one obtains N -acyclic \mathbb{I} -groupoids that are fully symmetric over \mathbb{I} .

11 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 [14] to yield a conceptual improvement and correction of the proposed constructions for groupoids from [15]. Due to its symmetry preserving generic character, the new presentation also supports the use of these groupoids in [16] 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 [6], which has meanwhile been treated comprehensively in [2].

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 \mathbf{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 [13, 14] and a precursor for special graphs in [9], we illustrate the key to the new generalisation in Proposition 11.1 below.
- (2) Natural reduced products with N -acyclic \mathbb{I} -groupoids yield finite branched N -acyclic coverings of hypergraphs $\mathbb{V} = (V, S)$ where a constraint pattern $\mathbb{I} = (S, E)$ is induced by the intersection graph of \mathbb{V} that encodes the intersection pattern between hyperedges in the given hypergraph (cf. [15]).
- (3) A new and more direct approach to finite branched N -acyclic coverings of hypergraphs $\mathbb{V} = (V, S)$ can be based on \mathbb{I} -products between a constraint graph $\mathbb{I} = (S, E)$ induced by the intersection graph of \mathbb{V} and suitable \mathbf{E} -groups that are not just N -acyclic but N -acyclic over \mathbb{I} ; cf. Proposition 11.3 below.

Of these fundamental applications, (2) has been explored in stages in [14, 15, 16]. Application (1) is new in its form that involves the new notion of N -acyclicity of groups over a constraint graph \mathbb{I} , rather than working with weak substructures of direct products with (unconstrained) N -acyclic Cayley groups. 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 9 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 \mathbf{E} -groups, the set of all edges $e = \{v, v'\} \in E$, and as a constraint graph \mathbb{I} the \mathbf{E} -graph $\mathbb{I} = (V, (\{e\})_{e \in E})$ (\mathbb{V} with individually labelled edges). For any \mathbf{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

$$\begin{aligned} \pi: \quad \hat{\mathbb{V}} &\longrightarrow \mathbb{V} \\ (v, g) &\longmapsto v \end{aligned}$$

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. Observation 6.6).

Proposition 11.1. *Let $\mathbb{V} = (V, E)$ be a connected finite simple graph, E associated with its edge set E as above and \mathbb{G} an \mathbf{E} -group that is compatible with*

the \mathbf{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 $\pi: (v, g) \mapsto v$.

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 $S \subseteq \mathcal{P}(V) \setminus \{\emptyset\}$ associate its *intersection graph* $\mathbb{I} = (S, E)$ where

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

If \mathbb{G} is an \mathbf{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} (cf. Definition 6.5) gives rise to a finite branched hypergraph covering of $\mathbb{V} = (V, S)$ by a hypergraph $\hat{\mathbb{V}} = (\hat{V}, \hat{S})$ as follows. Consider the disjoint union of $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ -tagged copies of the hyperedges of \mathbb{V} ,

$$\bigcup_{(s, g) \in \mathbb{I} \otimes \mathbb{C}\mathbb{G}} s \times \{(s, g)\} = \{(v, s, g) \in V \times S \times \mathbb{G} : v \in s\}$$

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

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

The induced equivalence is such that $(v, s, g) \approx (v', s', g')$ if, and only if, $v' = v \in s \cap s'$ and $g' = gh$ for some $h = [w]_{\mathbb{G}}$ with $w \in \alpha_v^*[\mathbb{I}, s, s']$, $\alpha_v := \{e = \{t, t'\} \in E : v \in t \cap t'\}$, which implies $g' \in \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_v, s; g]$. We write $[v, s, g]$ for the \approx -equivalence class of $(v, s, g) \in s \times \{(s, g)\}$. These equivalence classes form the vertices of $\hat{\mathbb{V}}$; its hyperedges the subsets induced by the $s \in S$, denoted as

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

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

$$\hat{\mathbb{V}} = (\hat{V}, \hat{S}) =: \mathbb{V} \otimes_{\mathbb{I}} \mathbb{C}\mathbb{G} \text{ where } \begin{cases} \hat{V} := \{[v, s, g] : v \in s, (s, g) \in \mathbb{I} \otimes \mathbb{C}\mathbb{G}\} \\ \hat{S} := \{[s, g] : (s, g) \in \mathbb{I} \otimes \mathbb{C}\mathbb{G}\} \end{cases}$$

with covering projection

$$\begin{aligned} \pi: \hat{\mathbb{V}} = (\hat{V}, \hat{S}) &\longrightarrow \mathbb{V} = (V, S) \\ [v, s, g] &\longmapsto v. \end{aligned}$$

Observation 11.2. Let $\hat{\mathbb{V}} := \mathbb{V} \otimes_{\mathbb{I}} \mathbb{C}\mathbb{G}$ be defined as above for an \mathbf{E} -group \mathbb{G} that is compatible with the intersection graph $\mathbb{I} = (S, E)$ of $\mathbb{V} = (V, S)$. Then hyperedges $[s_1, g_1]$ and $[s_2, g_2]$ of $\hat{\mathbb{V}}$ intersect precisely in vertices $[v, s_1, g_1] = [v, s_2, g_2]$ with $v \in s_1 \cap s_2$ where $g_2 \in \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_v, s_1; g_1]$, $\alpha_v := \{e = \{s, s'\} \in E : v \in s \cap s'\}$ (i.e. $g_2 = g_1[w]_{\mathbb{G}}$ for some w labelling an \mathbb{I} -walk from s_1 to s_2 in \mathbb{I} , or from (s_1, g_1) to (s_2, g_2) in $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$: $w \in \alpha_v^*[\mathbb{I}, s_1, s_2]$, cf. Definition 6.3 and Observation 6.6).

Proposition 11.3. Let (V, S) be a finite hypergraph, $\mathbb{I} = (S, E)$ its intersection graph, \mathbb{G} an \mathbf{E} -group that is compatible with \mathbb{I} . Then the hypergraph $\hat{\mathbb{V}} := \mathbb{V} \otimes_{\mathbb{I}} \mathbb{C}\mathbb{G}$ as defined above gives rise to a finite branched hypergraph covering $\pi: \hat{\mathbb{V}} \rightarrow \mathbb{V}$. The 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 [3, 4].

Proof. Consider the hypergraph $\mathbb{V} \otimes_{\mathbb{I}} \mathbb{C}\mathbb{G}$, as defined above, for an \mathbf{E} -group \mathbb{G} that is compatible with the intersection graph $\mathbb{I} = (S, E)$ of \mathbb{V} and N -acyclic for some $N \geq 3$. The covering property follows from Observation 11.2, and it remains to argue for N -acyclicity of $\hat{\mathbb{V}}$. We show that the Gaifman graph of $\hat{\mathbb{V}}$ cannot have chordless cycles of lengths n for $3 < n \leq N$ (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 hyperedges $[s_i, g_i] \in \hat{S}$ with $\hat{v}_i \in [s_i, g_i] \cap [s_{i+1}, g_{i+1}]$ witness the connectivity condition in this cycle. This implies that \hat{v}_i can be represented as $\hat{v}_i = [v_i, s_i, g_i] = [v_i, s_{i+1}, g_{i+1}]$ for some $v_i \in s_i \cap s_{i+1}$ and that $g_{i+1} = g_i[w_i]_{\mathbb{G}}$ for some $w_i \in \alpha_i^*[\mathbb{I}, s_i, s_{i+1}]$ where $\alpha_i = \{e = \{s, s'\} \in E : v_i \in s \cap s'\}$. 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. $g = g_i[w]_{\mathbb{G}}$ for some $w \in \alpha_{i,i-1}^*[\mathbb{I}, s_i, s]$ and $g := g_{i+1}[w']_{\mathbb{G}}$ for some $w' \in \alpha_{i,i+1}^*[\mathbb{I}, s_{i+1}, s]$ (the same s , due to compatibility of \mathbb{G} with \mathbb{I}). Then $\hat{v}_{i-1} = [v_{i-1}, s_{i-1}, g_{i-1}] = [v_{i-1}, s, g]$ 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, $\hat{v}_{i+1} = [v_{i+1}, s_{i+1}, g_{i+1}] = [v_{i+1}, s, g]$ 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 $m = \{\hat{v}_i : 1 \leq i \in \mathbb{Z}_n\}$ forms a clique of size $n \geq 3$ in the Gaifman graph of $\hat{\mathbb{V}} = (\hat{V}, \hat{S})$ such that every subset $m_i := m \setminus \{\hat{v}_i\}$ of size $n - 1$ is contained in some hyperedge (a minimal violation of conformality). For $i \in \mathbb{Z}_n$, let $[s_i, g_i] \in \hat{S}$ be a hyperedge that contains $m_i = m \setminus \{\hat{v}_i\} = \{\hat{v}_j : j \neq i\}$. So, for $j \neq i$, $\hat{v}_j \in [s_i, g_i]$ can be represented as $\hat{v}_j = [v_j, s_i, g_i] \in [s_i, g_i]$. Let $\gamma_i := \alpha_{v_i} = \{e = \{s, s'\} \in E : v_i \in s \cap s'\}$; and $\gamma_{i,i'} := \bigcap_{j \neq i, i'} \gamma_j$, so that \mathbb{I} -walks in $\gamma_{i,i'}$ preserve all \hat{v}_j for $j \neq i, i'$, in the sense that $[v_j, s_j, g_j] = [v_j, s, g_j]h$ for $h = [w]_{\mathbb{G}}$, $w \in \gamma_{i,i'}^*[\mathbb{I}, s_j, s]$. In particular, $g_{i+1} \in \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i; g_i]$ for $\alpha_i := \gamma_{i,i+1}$, as $\hat{v}_j \in [s_i, g_i] \cap [s_{i+1}, g_{i+1}]$ for $j \neq i, i+1$, cf. Observation 11.2. Consider then

$$(\mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i; g_i])_{i \in \mathbb{Z}_n} \quad \text{for } \alpha_i := \gamma_{i,i+1} = \bigcap_{j \neq i, i+1} \gamma_j$$

as a candidate for an \mathbb{I} -coset cycle. We show that if this were not an \mathbb{I} -coset cycle, then the whole of m would be contained in some hyperedge $[s, g] \in \hat{S}$. The connectivity condition on \mathbb{I} -coset cycles from Definition 6.8 is just that $g_{i+1} \in \mathbb{C}\mathbb{G}[\mathbb{I}, \alpha_i, s_i; g_i]$. The separation condition now is that

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

where $\alpha_{i,i-1} = \alpha_i \cap \alpha_{i-1} = \bigcap_{j \neq i} \alpha_j = \{e = \{s, s'\} \in E : m_i \subseteq s \cap s'\}$ and $\alpha_{i,i+1} = \alpha_i \cap \alpha_{i+1} = \bigcap_{j \neq i+1} \alpha_j = \{e = \{s, s'\} \in E : m_{i+1} \subseteq s \cap s'\}$.

Assume there were some g in the above intersection, i.e. $g = g_i h$ for some $h = [w]_{\mathbb{G}}$ with $w \in \alpha_{i,i-1}^*[\mathbb{I}, s_i, s]$ and $g = g_{i+1} h'$ for some $h' = [w']_{\mathbb{G}}$ with $w' \in \alpha_{i,i+1}^*[\mathbb{I}, s_{i+1}, s]$, for the same $s \in S$ (due to compatibility of \mathbb{G} with \mathbb{I}). We claim that this would imply $m \subseteq [s, g]$, contradicting the assumptions. For $j \neq i$, $\hat{v}_j \in [s_i, g_i]$ implies $\hat{v}_j = [v_j, s_i, g_i]$, which, together with $h = [w]_{\mathbb{G}}$, $w \in \alpha_{i,i-1}^*[\mathbb{I}, s_i, s]$ entails that $\hat{v}_j \in [s, g]$ for $j \neq i$. Similarly, for $j \neq i+1$, $\hat{v}_j \in [s_{i+1}, g_{i+1}]$ implies $\hat{v}_j = [v_j, g_{i+1}]$, which, together with $h' = [w']_{\mathbb{G}}$, $w' \in \alpha_{i,i+1}^*[\mathbb{I}, s_{i+1}, s]$ entails that $\hat{v}_j \in [s, g]$ for $j \neq i+1$. As \mathbb{G} is N -acyclic, $n > N$ follows. \square

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 [14] (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 [15], 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 [16] 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 [15], however, contained a serious flaw that was also carried in the arXive preprints. This mistake from [15] 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 [15] 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. Karl also provided extremely valuable critical comments on earlier versions of this paper. Those helped to eliminate some gaps and mistakes besides triggering a more specific treatment towards the Henkell–Rhodes conjecture in [2], which also recasts key concepts from the current paper in a more algebraic framework.

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] K. Auinger, J. Bitterlich, and M. Otto. Solution to the Henkell–Rhodes problem: finite F-inverse covers do exist. Preprint *arXiv:2208.03273*, 2022.
- [3] 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.
- [4] C. Berge. *Graphs and Hypergraphs*. North-Holland, 1973.
- [5] 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.
- [6] J. Bitterlich. Finite F-inverse covers do exist. Preprint *arXiv:1808.05161*, 2018.
- [7] A. Cayley. Desiderata and Suggestions: No. 1. The Theory of Groups. *American Journal of Mathematics*, vol. 1, pp. 50–52. 1878.
- [8] A. Cayley. Desiderata and Suggestions: No. 2. The Theory of Groups. *American Journal of Mathematics*, vol. 1, pp. 174–176. 1878.
- [9] 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.
- [10] R. Fox. Covering spaces with singularities. In *A Symposium in Honor of S. Lefschetz*, pp. 243–257. Princeton University Press, 1957.

- [11] R. Frucht. Herstellung von Graphen mit vorgegebener abstrakter Gruppe. *Compositio Mathematica*, 6, pp. 239–250, 1938.
- [12] M. Lawson. *Inverse Semigroups*. World Scientific, 1998.
- [13] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130, pp. 173–205, 2004.
- [14] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59 (1), 2012.
- [15] M. Otto. Groupoids, hypergraphs and symmetries in finite models. In *Proceedings of 28th Annual IEEE Symposium on Logic in Computer Science LICS'13*, 2013.
- [16] M. Otto. Amalgamation and symmetry: from local to global consistency in the finite. Technical report, 2018. See *arXiv*: 1709.00031.