Acyclicity in Finite Groups and Groupoids

Martin Otto
Department of Mathematics
Technische Universität Darmstadt

June 19, 2018

Abstract

We expound a simple 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 cyclic configurations formed 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 possess known applications in finite graph and hypergraph coverings that locally unfold cyclic configurations. An important new feature of the construction proposed here is that it reduces the hitherto considerably more complex construction for groupoids to a suitably adapted construction for groups, and even for groups with involutive generators, with the additional benefit of a more uniform approach across these settings.

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 [7] 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. A key ingredient 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.

On the other hand, interesting finite groups can be obtained as permutation groups, i.e. as subgroups of the full symmetric group of a finite set. These can be induced by graph-like extra structure on that set in other ways than just as the group of symmetries. Specific graph structures and carefully designed permutation group actions can thus give rise to finite groups with desirable algebraic or combinatorial properties suggested by various applications. A very nice example of this technique is a construction, due to Biggs [4], of finite groups over a given set of generators that avoid short cycles, i.e. 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 having 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 for 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 structure, 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 [10] for specific applications in finite graph coverings. These techniques were then substantially generalised to the setting of groupoids in [11], with specific applications towards hypergraph coverings (here necessarily branched, in a discrete analogue of classical terminology from [6]). The goal there and here again is a simple and generic combinatorial construction of groups and groupoids with strong acyclicity properties that control coset cycles rather than just generator cycles. One main technical point in the treatment of [11] over and above the ground work in [10] has to do with the difficulty to overcome the restriction

to involutive generators, which seem inadequate in a groupoidal setting. In the current, more comprehensive and more systematic extension of the original idea of [10] we propose a construction of highly acyclic finite groups with sets of involutive generators that yields stronger results – or stronger notions of acyclicity based on more general patterns than mere coset cycles. This allows us now to present a self-contained account in which even the much more involved construction of highly acyclic finite groupoids from [11] can be reduced to the new, enriched construction for groups with involutive generators. This results in a much simpler construction and a more 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 5.1 and 5.2.

Notation and conventions

In this paper we consider various kinds of graphs, some undirected, some directed, often also allowing loops (reflexive edges), and in Section 4 also multigraphs 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 one edge relation for each colour to encode edge-coloured graphs. In some instances, and especially in Section 4, 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 vertex of an 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).

For algebraic 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 $\mathsf{E}^* = \bigcup_{n\in\mathbb{N}} \mathsf{E}^n$. As is common in formal language theory, we write a typical word of length $n \in \mathbb{N}$ as $w = e_1 e_2 \cdots e_n \in \mathsf{E}^n$ (rather than e.g., in tuple notation, as $(e_1, e_2, \ldots, e_n)$), denoting its length as $n = |w|$. We also write, e.g. just $w_1 w_2$ for the *concatenation* of the words $w_1, w_2 \in \mathsf{E}^*$ (which is often denoted as $w_1 \cdot w_2$ with explicit

notation for the concatenation operation as a monoidal semigroup operation). The *empty word* $\lambda \in \mathsf{E}^*$, 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 $\mathsf{E}^*/\!\sim$, but we suppress this as an unnecessary distraction in our considerations.

2 Cayley & Biggs: the basic construction

The idea to associate groups with graphs, and vice versa, can be attributed to Arthur Cayley. The Cayley graph of an abstract group, w.r.t. to a chosen set of generators, encodes the algebraic structural information in 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 [7]. In particular, Cayley graphs are, by construction, not just regular but homogeneous in the sense of having a transitive automorphism group. On one hand, Cayley graphs thus 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 abstract groups with certain acyclicity properties from a permutation group action on suitably prepared graph structures is best illustrated by the basic example of a construction of regular graphs of high girth due to Biggs [4] and outlined in [1].

2.1 Biggs' construction

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

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

be the tree structure with vertex set

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

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

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

for $e \in \mathsf{E}$. By construction, each vertex $w \in V$ with $|w| < n$ is an interior vertex of $\mathbb{T}(\mathsf{E}, n)$ of degree $d = |\mathsf{E}|$, with precisely one R_e-neighbour for each $e \in \mathsf{E}$; the remaining vertices, viz. those $w \in V$ with $|w| = n$, are leaves of $\mathbb{T}(\mathsf{E}, n)$, each with an R_e-neighbour for a unique $e \in \mathsf{E}$ (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 $e \in \mathsf{E}$ we associate the permutation $\pi_e \in \mathrm{Sym}(V)$ that swaps any pair of vertices that are incident with a common e-coloured edge. This is the involutive permutation of V whose graph precisely is the matching R_e. The target of the construction is the group $\mathbb{G}(\mathsf{E}, n)$, which is the subgroup of $\mathrm{Sym}(V)$ generated by these involutions:

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

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

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

Its Cayley graph, w.r.t. the generators $(\pi_e)_{e \in \mathsf{E}}$, has as its vertex set the set of group elements $\rho \in G$ and edge relations

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

These edge relations are symmetric due to the involutive nature of the π_e in $\mathrm{Sym}(V)$, and they are irreflexive and pairwise disjoint since $\mathrm{id}_V \neq \pi_e \neq \pi_{e'}$ for $e \neq e'$, as can be seen most easily by their action as permutations on $\lambda \in V$. So this Cayley graph is a d-regular finite graph, whose automorphism group acts

transitively on the set of vertices. For the last claim consider the left action of the group on itself:

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

which clearly induces an automorphism of the Cayley graph (albeit not of the group, which is rigid once we label the generators). That the girth of the Cayley graph of $\mathbb{G}$ is at least $4n+2$ can be seen as follows. A reduced word $w \in \mathsf{E}^k$ of length $k \geqslant 1$ can be written as $w = e_1 u$. Let $v \in \mathsf{E}^n$ be a leaf of $\mathbb{T}(\mathsf{E}, n)$ whose reversal v^{-1} agrees with u (up to $\min(|u|, n)$). Applying the corresponding permutation $\pi_w = \pi_u \circ \circ \pi e_1$ to v, we see that the action of the permutations prescribed by the first (up to) $n+1$ letters of w takes that leaf step by step towards the root λ, the next n letters (if present) will take it step by step towards a different leaf, where the very next letter (if present) can have no effect so that it would take at least the action of another $2n$ letters after that to bring this vertex back to where we started. In other words, no reduced word of fewer than $n+1+n+1+2n = 4n+2$ letters can label a generator sequence that represents the neutral element of the group, which is the identity in $\mathrm{Sym}(V)$.

2.2 The basic format: groups with involutive generators

If we ask what is essential about this passage from an E-coloured graph (in the above case $\mathbb{T}(\mathsf{E}, n)$) to a group, 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 girth bound. If $\mathbb{T}(\mathsf{E}, n)$ were replaced, for instance, by the disjoint union of all E-coloured line graphs corresponding to reduced words $w \in \mathsf{E}^{2n}$, the above girth bound of $4n+2$ persists with essentially the same argument. (The trivial upper bound of $|V|!$ on the size of the resulting group and Cayley graph may well be affected, but we here aim for better control on cyclic configurations in finite groups and graphs, irrespective of sheer size.)

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 it involutive (its own inverse); this involution has precisely those vertices as fixed points at which the edge relation has loops, and its domain $\mathrm{dom}(R)$ and range $\mathrm{rng}(R)$ consists of the set of the vertices of degree 1. A *full matching* correspondingly 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-graphs]
For a set E, an E-*graph* is an undirected edge-coloured graph $\mathbb{H} = (V, (R_e)_{e \in \mathsf{E}})$ whose undirected edges are E-coloured in such a way that each R_e is a partial

matching over the vertex set V, i.e. each vertex $v \in V$ has degree at most 1 w.r.t. to R_e for each $e \in \mathsf{E}$. The E-graph $\mathbb{H} = (V, (R_e)_{e\in\mathsf{E}})$ is *complete* if each R_e is a full matching, i.e. if each vertex $v \in V$ has degree exactly 1 w.r.t. to R_e for each $e \in \mathsf{E}$. The *trivial completion* of an E-graph $\mathbb{H} = (V, (R_e)_{e\in\mathsf{E}})$ is the complete E-graph $\bar{\mathbb{H}} = (V, (\bar{R}_e)_{e\in\mathsf{E}})$ obtained by putting $\bar{R}_e := R_e \cup \{(v, v) \colon v \in V \setminus \mathrm{dom}(R_e)\}$.

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

For groups $\mathbb{G} = (G, \cdot, 1)$ (in multiplicative notation), an element $g \in G$ is an *involution* if $g = g^{-1}$. A subset $\mathsf{E} \subseteq G \setminus \{1\}$ 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-groups]
For a set E, an E-*group* is any group $\mathbb{G} = (G, \cdot, 1)$ that has $\mathsf{E} \subseteq G$ as a set of non-trivial involutive generators.[1]

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 $w \in \mathsf{E}^*$, so that

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

is a surjective homomorphism from the free monoid structure of E^* with concatenation and neutral element $\lambda \in \mathsf{E}^*$ 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* $e \in \mathsf{E}$ *(as represented by* reduced *words in* E^**) can be regarded as the* free involutive group *over* E*. E-groups as defined above are precisely the homomorphic images of this free involutive group over* E*.*

Definition 2.4. [sym($\mathbb{H}$)]
For an E-graph $\mathbb{H} = (V, (R_e)_{e\in\mathsf{E}})$ we let sym($\mathbb{H}$) be the subgroup of Sym(V) that is generated by the involutive permutations $\pi_e \colon V \to V$ induced by the full matchings of its trivial completion $\bar{\mathbb{H}} = (V, (\bar{R}_e)_{e\in\mathsf{E}})$. Provided the $(\pi_e)_{e\in\mathsf{E}}$ are pairwise distinct, we may regard sym($\mathbb{H}$) as an E-group by identifying $e \in \mathsf{E}$ with the generators π_e, for $e \in \mathsf{E}$.

Note that in terms of $\mathbb{H} = (V, (R_e)_{e\in\mathsf{E}})$ itself, π_e swaps the two vertices incident with any irreflexive edge of R_e and fixes all other vertices. The Biggs group $G(\mathsf{E}, n)$ as discussed above is sym($\mathbb{T}(\mathsf{E}, n)$). Recall our convention that permutations act from the right in terms of the group operation, which in terms

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

of composition is cast as $\pi_e\pi_{e'} = \pi_{e'} \circ \pi_e$. This extends to arbitrary words $w = e_1 \cdots e_n \in \mathsf{E}^*$ over E according to

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

which yields a surjective homomorphism from the free monoid structure of E^* with concatenation and neutral element $\lambda \in \mathsf{E}^*$ onto the group structure of $\mathrm{sym}(\mathbb{H})$ with composition and neutral element $\pi_\lambda = \mathrm{id}_V$. Factorisation w.r.t. the identities $e = e^{-1}$ as in Observation 2.3, which are reflected in $\mathrm{sym}(\mathbb{H})$, turns this into a surjective group homomorphism from the free involutive group over E onto $\mathrm{sym}(\mathbb{H})$.

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

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

of colour e, for each generator $e \in \mathsf{E}$.

Note that a 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 $e \in G$, $\pi_e \colon g \mapsto ge$. It is easy to check that, as an abstract group with generators $e \in \mathsf{E}$, $\mathbb{G}$ is isomorphic to the subgroup of the full symmetric group $\mathrm{Sym}(G)$ over the vertex set G generated by these permutations π_e. In particular, in the case of a group $\mathbb{G} = (G, \cdot, 1)$ that admits a set of involutive generators $\mathsf{E} \subseteq G$, the associated Cayley graph $\mathbb{C}\mathbb{G} = \mathrm{Cayley}(\mathbb{G}, \mathsf{E})$ is a complete E-graph in the sense of Definition 2.1, and

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

For a not necessarily complete E-graph $\mathbb{H}$ with trivial completion $\bar{\mathbb{H}}$, the Cayley graph of $\mathrm{sym}(\mathbb{H}) = \mathrm{sym}(\bar{\mathbb{H}})$ is another, non-trivial complete E-graph that can be associated with $\mathbb{H}$. Moreover, by the above, the map

$$\mathbb{H} = (V, (R_e)_{e\in\mathsf{E}}) \longmapsto \mathrm{Cayley}\big(\mathrm{sym}(\bar{\mathbb{H}}), \{\pi_e \colon e \in \mathsf{E}\}\big)$$

is a projection from the class of all E-graphs onto the subclass of complete E-graphs.

In the following it will be convenient, and without risk of confusion, to identify the generators $e \in \mathsf{E}$ of a group $\mathbb{G}$ with the maps $\pi_e \colon g \mapsto ge$ in $\mathbb{G}$ or in its Cayley graph, and similarly to identify the family of generators $(\pi_e)_{e\in\mathsf{E}}$ of $\mathrm{sym}(\mathbb{H})$ with a subset $\mathsf{E} \subseteq \mathrm{sym}(\mathbb{H})$ by writing just e instead of π_e in this context, too. In this sense the above projection operation on E-graphs $\mathbb{H}$, for instance, simplifies to $\mathbb{H} \mapsto \mathrm{Cayley}(\mathrm{sym}(\bar{\mathbb{H}}), \mathsf{E})$.

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

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

The Cayley graph $\mathbb{CG}[\alpha]$ of $\mathbb{G}[\alpha]$, correspondingly is regarded as an α-graph, i.e. as a weak subgraph of the Cayley graph of $\mathbb{G}$ (more specifically it is the $(R_e)_{e\in\alpha}$-reduct of an induced subgraph on $G[\alpha] \subseteq G$).

Definition 2.7. [generated subgraphs]
For a subset $\alpha \subseteq \mathsf{E}$ and an E-graph $\mathbb{H} = (V, (R_e)_{e\in\mathsf{E}})$, an *$\alpha$-walk* of length n from v to v' is a sequence $v_0, e_1, v_1, e_2, \ldots, e_n, v_n$ of vertices, where $v_i \in V$, $v = v_0$, $v_n = v'$, and edge labels $e_i \in \alpha$ such that $(v_i, v_{i+1}) \in R_{e_{i+1}}$ for $i < n$.

The *α-component* of $v \in V$ consists of those vertices v' that are linked to v by α-walks. We write $\alpha[v] \subseteq V$ for this set of vertices and $\mathbb{H}[\alpha; v]$ for the α-graph induced by $\mathbb{H}$ on $\alpha[v]$.

As an $(R_e)_{e\in\alpha}$-reduct of the E-graph $\mathbb{H}$, this again is in general a weak rather than an induced subgraph. Clearly the Cayley graph of $\mathbb{G}[\alpha]$ arises as $\mathbb{H}[\alpha; 1]$ for $\mathbb{H} = \mathrm{Cayley}(\mathbb{G}, \mathsf{E}) = \mathbb{CG}$, so that the same structure $\mathbb{CG}[\alpha] = \mathbb{CG}[\alpha; 1]$ arises in two different manners.

It is useful to note that, if $v = v_0, e_1, v_1, e_2, \ldots, e_n, v_n = v'$ is an α-walk from v to v' in $\mathbb{H}$ s.t. $w = e_1 \cdots e_n$ traces the labels along this walk, then $v' = \pi_w(v) = [w]_{\mathbb{H}}(v)$.

Definition 2.8. [compatibility]
An E-group $\mathbb{G}$ is *compatible* with the E-graph $\mathbb{H} = (V, (R_e)_{e\in\mathsf{E}})$ if, for all $w \in \mathsf{E}^*$,

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

Clearly any E-group $\mathbb{G}$ is compatible with its Cayley graph. Note that $\mathrm{sym}(\mathbb{H})$ is compatible with $\mathbb{H}$ by construction. More specifically, for $\mathbb{G} = \mathrm{sym}(\mathbb{H})$, $\mathbb{G}$ is compatible with every connected component of $\mathbb{H}$, and for $\alpha \subseteq \mathsf{E}$, the generated subgroup $\mathbb{G}[\alpha] \subseteq \mathbb{G}$ is compatible, as an α-group, with every α-component $\mathbb{H}[\alpha; v]$ of $\mathbb{H}$. Generally, compatibility of $\mathbb{G}$ with $\mathbb{H}$ precisely guarantees that $\mathrm{sym}(\mathbb{H})$ is a homomorphic image of $\mathbb{G}$ under the natural mapping $[w]_{\mathbb{G}} \mapsto [w]_{\mathbb{H}}$. It is also worth noting that compatibility of the E-group $\mathbb{G}$ with several given E-graphs is equivalent to compatibility with their disjoint union.

The following definition involves a simple criterion on the relationship between α-components for different subsets $\alpha \subseteq \mathsf{E}$; it will later also feature as the first non-trivial level of coset-acyclicity in E-groups, cf. Definition 3.2.

Definition 2.9. [simple connectivity/2-acyclicity]
An E-graph is called *simply connected* if, for all $\alpha_1, \alpha_2 \subseteq \mathsf{E}$ and vertices v

$$\alpha_1[v] \cap \alpha_2[v] = (\alpha_1 \cap \alpha_2)[v].$$

An E-group $\mathbb{G}$ is simply connected (2-acyclic) if its Cayley graph $\mathbb{CG}$ is.

Example 2.10. An ordinary cycle with E-labelled edges that can be split into two connected pieces labelled by disjoint sets of generators violates simple connectivity. It similarly violates simple connectivity if any one of its contractions to some subset $\alpha \subseteq \mathsf{E}$ (obtained by contraction of e-edges for $e \notin \alpha$) splits in this manner. Conversely, it can be checked that a plain generator cycle is simply connected (2-acyclic) if none of its α-contractions splits into two disjointly labelled connected pieces. In particular, any periodic cyclic labelling of a cycle of the form w^n for any $w \in \mathsf{E}^*$ and $n \geqslant 2$ is simply connected, while a labelling of a cycle by $w_1 w_2$ with $w_i \in \alpha_i^* \setminus \{\lambda\}$, $\alpha_1 \cap \alpha_2 = \emptyset$ fails to be simply connected.

2.3 Amalgams of Cayley graphs

Definition 2.11. [α-similar to $\mathbb{G}$]
Let $\mathbb{G}$ be an E-group, $\alpha \subseteq \mathsf{E}$. An E-graph $\mathbb{H} = (V, (R_e)_{e \in \mathsf{E}})$ is *α-similar to $\mathbb{G}$ at $v \in V$*, denoted $\mathbb{H}, v \sim_\alpha \mathbb{G}$, if the α-generated subgraph $\mathbb{H}[\alpha; v]$ of $\mathbb{H}$ is isomorphic to a weak subgraph of the Cayley graph of $\mathbb{CG}[\alpha]$.

We note that an isomorphism of $\mathbb{H}[\alpha; v]$ with a weak subgraph of $\mathbb{CG}[\alpha] \subseteq \mathbb{CG}$ is uniquely determined by the choice of a group element $g \in G[\alpha]$ to be associated with v: if the α-walk $v = v_0, e_1, v_1, e_2, \ldots, e_n, v_n = v'$ links $v' \in \alpha[v]$ to v in $\mathbb{H}$, then v' must be mapped to $g \cdot [w]_{\mathbb{G}}$ by any isomorphism that maps v to g. Moreover the choice of $g \in G[\alpha]$ is completely free, due to homogeneity of $\mathbb{CG}$: $\mathbb{CG}[\alpha; g] \simeq \mathbb{CG}[\alpha; 1] = \mathbb{CG}[\alpha]$ for any $g \in G[\alpha]$.

If $\mathbb{H} = (V, (R_e)_{e \in \mathsf{E}})$ is α-similar to $\mathbb{G}$ at v, then an isomorphic copy of $\mathbb{CG}[\alpha], 1$ can be amalgamated with $\mathbb{H}$ in $v \in V$ in a unique manner by identifying v with with $1 \in G[\alpha]$.

Definition 2.12. [amalgamation]
Let $\mathbb{H} = (V, (R_e)_{e \in \mathsf{E}})$ be α-similar to $\mathbb{G}$ at $v \in V$. Then the *amalgamation* of $\mathbb{H}, v$ with $\mathbb{CG}[\alpha], 1$,

$$(\mathbb{H}, v) \oplus (\mathbb{CG}[\alpha], 1)$$

is the E-graph obtained by superposing the structure of $\mathbb{CG}[\alpha]$ as a complete α-graph with $\mathbb{H}[\alpha; v]$ via the unique isomorphism that identifies of $1 \in G[\alpha]$ with $v \in V$.

It is useful to note that

$$(\mathbb{H}, v') \oplus (\mathbb{CG}[\alpha], 1) \;\simeq\; (\mathbb{H}, v) \oplus (\mathbb{CG}[\alpha], 1)$$

for any $v' \in \alpha[v]$ in $\mathbb{H}$. An isomorphism of E-graphs is obtained as the combination of the identity on the (image of the) vertex set of V of $\mathbb{H}$ in $(\mathbb{H}, v') \oplus (\mathbb{CG}[\alpha], 1)$ with the inner automorphism of $\mathbb{CG}[\alpha]$ that maps $1 \in G[\alpha]$ to the element $g' \in G[\alpha]$ that gets identified with v' in $(\mathbb{H}, v) \oplus (\mathbb{CG}[\alpha], 1)$.

The following shows that simple connectivity (2-acyclicity) of E-graphs is *not* generally preserved in amalgamations. This is why we shall impose a stronger condition than just local similarity in the definition of amalgamation chains below (cf. Definition 2.14 and Lemma 2.15).

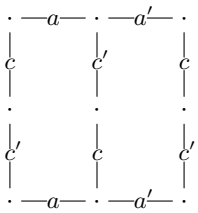

Figure 1: Amalgamated cycles as in Example 2.13

Example 2.13. Let a, a', c, c' be pairwise distinct, $\alpha = \{a, c, c'\}$, $\alpha' = \{a', c, c'\}$ and consider the α-labelled cycle $ac'cac'c$ and its α'-variant $a'c'ca'c'c$. These two cycles are 2-acyclic as can be verified directly or by reference to Example 2.10. If we join them by amalgamating two of their $c'c$-segments as in Figure 1, the amalgam consists of an $(\alpha \cup \alpha')$-cycle labelled $ac'caa'cc'a'$ with a single chord labelled $c'c$. This cycle splits into components w.r.t. $\gamma = \{a, a', c\}$ and $\gamma' = \{a, a', c'\}$ that violate simple connectivity.

Definition 2.14. [amalgamation chain]
Let $\mathbb{G}$ be an E-group $\mathbb{G}$ and $(\mathbb{G}[\alpha_i], g_i)_{1 \leqslant i \leqslant n}$ a sequence of subgroups generated by subsets $\alpha_i \subseteq \mathsf{E}$ together with distinguished elements $g_i \in \mathbb{G}[\alpha_i]$. The *amalgamation chains* $\mathbb{H}_k := \oplus_{i=1}^{k}(\mathbb{CG}[\alpha_i], g_i)$ are conditionally defined by induction on $1 \leqslant k \leqslant n$, together with distinguished vertices v_k in $\mathbb{H}_k$ according to:

(i) $\mathbb{H}_1, v_1 := \mathbb{CG}[\alpha_1], g_1$ for $k = 1$ (unconditionally);

(ii) for $k < n$, and under the condition that $\alpha_{k+1}[v_k] \subseteq \alpha_k[v_k]$ in $\mathbb{H}_k$, let

$$\mathbb{H}_{k+1} := (\mathbb{H}_k, v_k) \oplus (\mathbb{CG}[\alpha_{k+1}], 1)$$

and v_{k+1} the vertex corresponding to g_{k+1} in the amalgamated $\mathbb{CG}[\alpha_{k+1}]$.

In particular, the *amalgamation chain* $\mathbb{H} = \bigoplus_{i=1}^{n}(\mathbb{CG}[\alpha_i], g_i)$ of *length* n is defined as the final stage $\mathbb{H} := \mathbb{H}_n$ if, and only if, the required inclusion of α-components holds in all intermediate steps.

Well-definedness of $\mathbb{H}_{k+1}$ follows from the fact that the inclusion condition on α-components at v_k in $\mathbb{H}_k$ implies that $\mathbb{H}_k$ is α_{k+1}-similar to $\mathbb{G}$ at v_k, viz. that $\mathbb{H}_k[\alpha_{k+1}; v_k] \simeq \mathbb{H}_k[\alpha_{k+1} \cap \alpha_k; v_k] \simeq \mathbb{CG}[\alpha_{k+1} \cap \alpha_k] \subseteq \mathbb{CG}[\alpha_k]$ by construction.

The following lemma will be crucial towards certain compatibility arguments to be considered in main constructions in the next section.

Lemma 2.15. *Provided the constituents $\mathbb{CG}[\alpha_i]$ are 2-acyclic, the β-components of vertices in an amalgamation chain $\mathbb{H} = \bigoplus_{i=1}^{n}(\mathbb{CG}[\alpha_i], g_i)$ arise as amalgamation chains of lengths up to n based on constituents $\mathbb{CG}[\alpha_i \cap \beta]$.*

Proof. Consider a single amalgamation step from $\mathbb{H}_k$ to $\mathbb{H}_{k+1}$ in the amalgamation chain $\mathbb{H} = \bigoplus_{i=1}^{n}(\mathbb{CG}[\alpha_i], g_i)$. Assume that $\mathbb{G}[\alpha]$ for $\alpha := \alpha_k$ is 2-acyclic and that v_k in $\mathbb{H}_k$ itself is an element of the β-component in question (cf. discussion

after Definition 2.14 regarding the flexibility w.r.t. the actual choice of v in its α_{k+1}-component).

We choose to regard $\mathbb{CG}[\alpha_k] \simeq \mathbb{H}[\alpha, v]$ as an actual induced substructure $\mathbb{CG}[\alpha_k] = \mathbb{H}[\alpha; v] \subseteq \mathbb{H}_k$ of $\mathbb{H}_k$ so that $v := v_k = g_k$, and assume further that $\mathbb{H}_k, v_k$ satisfies the necessary condition for continuation of the amalgamation chain from $\mathbb{H}_k$ to $\mathbb{H}_{k+1}$: $\alpha_{k+1}[v] \subseteq \alpha[v]$ in $\mathbb{H}_k$. The overlap between the copy of $\mathbb{CG}[\alpha_{k+1}]$ that is amalgamated with $\mathbb{H}_k$ to obtain $\mathbb{H}_{k+1}$ is precisely $\mathbb{H}_k[\alpha_{k+1}; v]$. It follows that the β-component of v in $\mathbb{H}_{k+1}$ consists of the union of $\mathbb{H}_k[\beta, v]$ with the β-components in the copy of $\mathbb{CG}[\alpha_{k+1}]$ of all those elements that are both in the β- and in the α_{k+1}-component of v in $\mathbb{CG}[\alpha]$. These are the β-components in the amalgamated copy of $\mathbb{CG}[\alpha_{k+1}]$ of elements identified with

$$\beta[v] \cap \alpha_{k+1}[v] \text{ in } \mathbb{CG}[\alpha].$$

By 2-acyclicity of $\mathbb{CG}[\alpha]$, this intersection of components is just $(\alpha_{k+1} \cap \beta)[v]$. It follows that

$$\mathbb{H}_{k+1}[\beta, v] \simeq (\mathbb{H}_k[\beta, v], v) \oplus (\mathbb{CG}[\alpha_{k+1} \cap \beta], 1).$$

To use this argument in an inductive proof of the claim of the lemma, which speaks about any one β-component in a given amalgamation chain $\mathbb{H} = \bigoplus_{i=1}^{n}(\mathbb{CG}[\alpha_i], g_i)$ over 2-acyclic $\mathbb{CG}[\alpha_i]$, we initialise it for $\mathbb{H}_k$ where k is the least k such that the β-component in question intersects $\mathbb{H}_k$, and then continue along amalgamation steps in the chain for as long as the β-component in question does not die out. □

3 Acyclicity properties

3.1 Coset cycles vs. generator cycles

The immediate notion of a cycle in an E-group $\mathbb{G}$ is based on the graph theoretic notion of cycles in its Cayley graph $\mathbb{CG}$. Such a *generator cycle* of length $n \geqslant 3$ is traced out by a cyclically indexed tuple of group elements $(g_i)_{i \in \mathbb{Z}_n}$ where $g_{i+1} = g_i e_i$ for generators $e_i \in \mathsf{E}$ with $e_{i+1} \neq e_i$. Algebraically, therefore, this cycle is induced by a tuple $(e_i)_{i \in \mathbb{Z}_n}$ of generators (corresponding to a reduced cyclic word in E^n) with $e_{i+1} \neq e_i$ and

$$\textstyle\prod_{i=1}^{n} e_i = 1.$$

We note that this equality is invariant under cyclic permutations since the e_i are involutions. As discussed in Section 2, Biggs' construction yields finite E-groups of any desired finite degree of acyclicity w.r.t. such generator cycles, i.e. finite E-groups and Cayley graphs of arbitrarily large girth. We here want to focus on much stronger forms of acyclicity based on forbidding cyclic configurations formed by more complex building blocks – in the first instance, by cosets. I.e., we control short *coset cycles* rather than just short generator cycles. It will

be apparent from the definitions below that generator cycles are very special kinds of coset cycles.

A (left) *coset* in an E-group $\mathbb{G}$ is any subset of the form $g\mathbb{G}_0 = \{gh\colon h \in G_0\}$, where $\mathbb{G}_0 \subseteq \mathbb{G}$ is a subgroup of $\mathbb{G}$ and $g \in G$ any element of the group $\mathbb{G}$; in the following we are interested in cosets formed by generated subgroups, i.e. cosets of the form

$$g\mathbb{G}[\alpha] = \{gh\colon h \in \mathbb{G}[\alpha]\} = \{g \cdot [w]_{\mathbb{G}}\colon w \in \alpha^*\} \subseteq G,$$

for subsets $\alpha \subseteq \mathsf{E}$. As a set of vertices in the Cayley graph $\mathbb{C}\mathbb{G}$, the coset $g\mathbb{G}[\alpha]$ arises as the α-component of the vertex g. A *pointed coset* is a coset $g\mathbb{G}[\alpha]$ together with a distinguished element g that gives rise to it.

Definition 3.1. [coset cycles]
For $n \geqslant 2$, a *coset cycle of length* n in an E-group $\mathbb{G}$ is a cyclically indexed tuple of pointed cosets $(g_i\mathbb{G}[\alpha_i], g_i)_{i\in\mathbb{Z}_n}$ such that, for all i,

(i) $g_{i+1} \in g_i\mathbb{G}[\alpha_i]$;

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

Condition (i) just says that consecutive cosets along the cycle do overlap (in the named group elements); condition (ii) requires consecutive overlaps to be locally separate in the sense that the predecessor and successor cosets of the link $g_i\mathbb{G}[\alpha_i]$ do *not intersect within* $g_i\mathbb{G}[\alpha_i]$.

Looking at the group elements $h_i := g_i^{-1}g_{i+1}$ that link the consecutive distinguished elements in which the cosets overlap, we see that by (i), $h_i \in \mathbb{G}[\alpha_i]$ and that

$$\textstyle\prod_{i=1}^{n} h_i = 1.$$

3.2 Coset acyclicity

Definition 3.2. [N-acyclicity]
For $N \geqslant 2$, an E-group $\mathbb{G}$ is *N-coset-acyclic* (*N-acyclic* for short) if it admits no coset cycles of lengths up to N.

Note that 2-acyclicity of $\mathbb{G}$ coincides with the previously defined notion of *simple connectivity* of $\mathbb{C}\mathbb{G}$ from Definition 2.9 – which albeit somewhat degenerate is by no means trivial.

The following remark makes a connection with acyclicity notions from classical hypergraph theory. It thereby illustrates the relevance of coset acyclicity from a different angle, but we shall here not pursue that connection further.

Remark 3.3. *With an E-group $\mathbb{G}$ or its Cayley graph $\mathbb{C}\mathbb{G}$ we may associate its* dual hypergraph $\mathbb{D}\mathbb{G}$ *whose vertex set is the set of all cosets induced by generated subgroups, clustered into hyperedges by the group elements they share:*

$$\mathbb{D}\mathbb{G} := (V, S) \;\; \textit{for} \;\; V := \{g\mathbb{G}[\alpha]\colon \alpha \subseteq \mathsf{E}, g \in G\}$$
$$S := \{[\![g]\!]\colon g \in G\} \;\; \textit{where} \;\; [\![g]\!] = \{g\mathbb{G}[\alpha]\colon \alpha \subseteq \mathsf{E}\}.$$

The Gaifman graph *of a hypergraph has the same vertex set and edges that turn every hyperedge into a clique. In classical hypergraph terminology, cf. [3, 2], a hypergraph is* acyclic *(also referred to as α-acyclicity, and equivalent to tree-decomposability) if its Gaifman graph is conformal and chordal. For 2-acyclic Cayley groups (cf. simple connectivity), shortest coset cycles in $\mathbb{G}$ of lengths greater then 3 can be directly related to chordless cycles in the Gaifman graph of $\mathbb{DG}$, while coset cycles of length 3 correspond to triangles in the Gaifman graph of $\mathbb{DG}$ that are not contained in any single hyperedge of $\mathbb{DG}$. It is then not hard to see that for Cayley groups that are at least 2-acyclic, the following are equivalent, for every $N \geqslant 3$:*

(i) N-acyclicity of $\mathbb{G}$: no coset cycles of length up to N;

(ii) N-acyclicity of the dual hypergraph $\mathbb{DG}$: every induced sub-hypergraph on up to N vertices is acyclic as a hypergraph (chordal and conformal).

The following construction is essentially presented in [11], but we here use a variation of the proof that will allow for new extensions further below.

Proposition 3.4. *For every finite set E and every $N \geqslant 2$ one can construct finite E-groups $\mathbb{G}$ that are compatible with any amalgamation chains of lengths up to N of subgroups $\mathbb{CG}[\alpha_i]$ generated by any subsets $\alpha_i \subseteq \mathsf{E}$. In particular, such $\mathbb{G}$ is guaranteed to be N-acyclic. Moreover, the E-groups obtained in the proposed construction are fully symmetric w.r.t. permutations of the set E of their involutive generators.*

In addition, the finite E-group can be required to be compatible with a given finite E-graph $\mathbb{H}$, and symmetric w.r.t. all those permutations of the set E that are also symmetries of the given E-graph $\mathbb{H}$.

Proof. The desired N-acyclic E-group $\mathbb{G}$ is obtained in an inductive process that guarantees compatibility within amalgamation chains of subgroups of $\mathbb{G}[\alpha]$ for subsets $\alpha \subseteq \mathsf{E}$ of growing sizes $|\alpha|$. More specifically, we inductively construct a sequence of E-groups $\mathbb{G}_n$ for $n \leqslant |\mathsf{E}|$ such that

(i) the sequence of the $\mathbb{G}_n$ is conservative w.r.t. generated subgroups in the sense that $\mathbb{G}_n[\alpha] \simeq \mathbb{G}_m[\alpha]$ whenever $n, m \geqslant |\alpha|$;

(ii) for any $\alpha \subseteq \mathsf{E}$ and all $n \geqslant |\alpha|$, the generated subgroup $\mathbb{G}_n[\alpha] \subseteq \mathbb{G}_n$ is compatible with amalgamation chains of lengths up to N of generated subgroups, in the sense that $\mathbb{G}_n[\alpha]$ (as an α-group) is compatible with any amalgamation chain of the form $\oplus_{i=1}^{k}(\mathbb{G}_n[\alpha \cap \alpha_i], g_i)$ for $k \leqslant N$.

By (ii), $\mathbb{G} := \mathbb{G}_{|\mathsf{E}|}$ is as desired: we note that compatibility with amalgamated subgroups as in condition (ii) implies that, for $|\alpha| \leqslant n$, $\mathbb{G}_n[\alpha]$ is N-acyclic as an α-group. Towards this claim consider any candidate coset cycle of length $2 \leqslant k \leqslant N$ in the α-group $\mathbb{G} := \mathbb{G}_n[\alpha]$:

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

where $\alpha_i \subseteq \alpha$, with $h_i := g_i^{-1}g_{i+1} \in \mathbb{G}[\alpha_i]$. By (ii), $\mathbb{G} = \mathbb{G}_n[\alpha]$ is compatible with the amalgamation chain

$$(**) \qquad \mathbb{H} = \oplus_{i=1}^{k}(\mathbb{CG}_n[\alpha_i], h_i)$$

for $h_i := g_i^{-1}g_{i+1}$. Comapare conditions (i) and (ii) for coset cycles in Definition 3.1. Condition (i) there implies that $h_i \in \mathbb{G}_n[\alpha_i]$; and condition (ii) implies that $(**)$ satisfies the inclusion requirement between components for amalgamation chains (compare Definition 2.14): $g_j\mathbb{G}[\alpha_j \cap \alpha_{j-1}] \cap g_{j+1}\mathbb{G}[\alpha_j \cap \alpha_{j+1}] = \emptyset$ implies that $\alpha_{j+1}[h_j] \subseteq \alpha_j[h_j]$ in the $\mathbb{G}[\alpha_j]$-summand of this chain. By induction then also $\alpha_{j+1}[h_j] \subseteq \alpha_j[h_j]$ in $\oplus_{i=1}^{j}(\mathbb{G}[\alpha_i], h_i)$, because the $\mathbb{G}[\alpha_j]$ overlaps with $\oplus_{i=1}^{j-1}(\mathbb{G}[\alpha_i], h_i)$ only in $\alpha_j[1]$. But then compatibility of $\mathbb{G}$ with $\mathbb{H}$ implies that $(*)$ cannot be a coset cycle: tracing the vertex 1 in the copy of $\mathbb{C}\mathbb{G}_n[\alpha_1]$ in $\mathbb{H}$ according to $(**)$ under the operation of the $\pi_i = \pi_{w_i} = [w_i]_{\mathbb{H}}$ for generator words $w_i \in \alpha_i^*$ that generate the h_i within $\mathbb{G}[\alpha_i]$, for $i = 1, \ldots, k$, we see that it gets mapped to the vertex g_k in the copy of $\mathbb{C}\mathbb{G}_n[\alpha_k]$ in $\mathbb{H}$, which is not the operation of id_V on $\mathbb{H}$ as required if $(*)$ were a coset cycle. So $\prod h_i \neq 1$ in $\mathbb{G} = \mathbb{G}_n[\alpha]$, as $\mathbb{G}$ is compatible with $\mathbb{H}$ by construction.

The inductive construction. Let $\mathbb{G}_1$ be $\mathbb{G}_1 := \mathbb{Z}_2^{\mathsf{E}}$, the additive group of an $|\mathsf{E}|$-dimensional $\mathbb{Z}_2$-vector space, in which the standard basis vectors play the rôle of the involutive generators. This group is also obtained as $\mathbb{G}_1 = \mathrm{sym}(\mathbb{H}_0)$ for the E-graph $\mathbb{H}_0$ consisting of a disjoint union of individual e-edges, one for each $e \in \mathsf{E}$. Inductively we obtain $\mathbb{G}_{n+1}$ from $\mathbb{G}_n$ as

$$\mathbb{G}_{n+1} := \mathrm{sym}(\mathbb{H}_n),$$

where $\mathbb{H}_n$ is the E-graph obtained as the disjoint union of

- the Cayley graph $\mathbb{C}\mathbb{G}_n$ of $\mathbb{G}_n$;
- all E-graphs $\oplus_{i=1}^{k}(\mathbb{C}\mathbb{G}_n[\alpha_i], g_i)$ that arise as amalgamation chains of length $k \leqslant N$ of Cayley graphs of subgroups $\mathbb{G}_n[\alpha_i] \subseteq \mathbb{G}_n$ based on generator subsets $\alpha_i \subseteq \mathsf{E}$ of sizes $|\alpha_i| \leqslant n$.

It remains to argue that the $\mathbb{G}_n$ thus defined satisfy conditions (i) and (ii) above. Condition (i) is guaranteed for $|\alpha| \leqslant 1$, since $\mathbb{G}[\emptyset] \simeq \{1\}$ and $\mathbb{G}[\{e\}] \simeq \mathbb{Z}_2$ in any E-group $\mathbb{G}$. Also condition (ii) is vacuously true even of any E-group $\mathbb{G}$, since the subgroups $\mathbb{G}[\alpha]$ for $|\alpha| \leqslant 1$ are either the trivial group $(\{1\}, \cdot, 1)$ or isomorphic to $(\mathbb{Z}_2, +, 0)$, neither of which admits any non-trivial amalgamation chains.

For the step from $\mathbb{G}_n$ to $\mathbb{G}_{n+1}$ we first observe that condition (i) is preserved, i.e. that any subgroup $\mathbb{G}_n[\alpha]$ for $|\alpha| \leqslant n$ is isomorphic to its sibling $\mathbb{G}_{n+1}[\alpha]$ in $\mathbb{G}_{n+1}$ for the following reason. As $\mathbb{G}_{n+1} = \mathrm{sym}(\mathbb{H}_n)$, $\mathbb{G}_{n+1}[\alpha]$ is the same as $\mathrm{sym}(\mathbb{H}_n^{\alpha})$ where $\mathbb{H}_n^{\alpha}$ is the disjoint union of all α-components of vertices in $\mathbb{H}_n$. As the $\mathbb{G}_n[\alpha_i]$ that contribute to amalgamation chains in $\mathbb{H}_n$ are all 2-acyclic, Lemma 2.15 implies that the α-components of those amalgamation chains are themselves amalgamation chains based on $\mathbb{G}_n[\alpha_i \cap \alpha]$; these therefore are amalgamation chains of length 1 and this of the form $\mathbb{G}_n[\alpha_i]$ for $|\alpha_i| \leqslant n$ or, if of lengths greater than 1, based on several $\mathbb{G}_n[\alpha_i \cap \alpha]$ for which $|\alpha_i \cap \alpha| < n$. It follows that already $\mathbb{G}_n$ is compatible with $\mathbb{H}_n^{\alpha}$, whence $\mathbb{G}_{n+1}[\alpha] \simeq \mathbb{G}_n[\alpha]$.

That $\mathbb{G}_{n+1}$ satisfies condition (ii) means that $\mathbb{G}_{n+1}[\alpha]$ is compatible with any amalgamation chain of generated subgroups of length up to N for $|\alpha| \leqslant n+1$. For $|\alpha| \leqslant n$, this is directly inherited from $\mathbb{G}_n$, due to (i) and since $\mathbb{G}_{n+1}$ is

compatible with $\mathbb{G}_n$. For $|\alpha| = n+1$, consider any amalgamation chain

$$(\dagger) \qquad \mathbb{H} = \oplus_{i=1}^{k}(\mathbb{C}\mathbb{G}_{n+1}[\alpha_i], g_i)$$

of length $k \leqslant N$ in the α-group $\mathbb{G} := \mathbb{G}_{n+1}[\alpha]$, for $\alpha_i \subseteq \alpha$. If $\alpha_i = \alpha$ for some i, then by the criteria of Definition 2.14 this chain is isomorphic to the single summand $\mathbb{C}\mathbb{G}_{n+1}[\alpha]$, and $\mathbb{G}_{n+1}[\alpha]$ is trivially compatible with its own Cayley graph. In all non-trivial amalgamation chains of the form (†), therefore, $\alpha_i \subsetneq \alpha$ so that $|\alpha_i| \leqslant n$, and, by condition (i), $\mathbb{H} \simeq \oplus_{i=1}^{k}(\mathbb{C}\mathbb{G}_n[\alpha_i], g_i)$, which is one of the amalgamation chains in $\mathbb{H}_n$ that $\mathbb{G}_{n+1}$ is compatible with by construction.

It is obvious that the construction of the $\mathbb{G}_n$ is symmetric w.r.t. any permutation of the generator set E. If we also want the $\mathbb{G}_n$ to be compatible with a given finite E-graph $\mathbb{H}$, we may start from $\mathbb{G}_1 := \mathrm{sym}(\mathbb{Z}_2^{\mathsf{E}} \dot\cup \mathbb{H})$ instead of the default choice above; in that case all $\mathbb{G}_n$ are still symmetric w.r.t. any permutation of E that is also a symmetry of $\mathbb{H}$. □

3.3 Variants of coset acyclicity

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 $\alpha^* \subseteq \mathsf{E}^*$ under the natural homomorphism

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

that associates a group element with any (reduced) word over E. This association naturally extends to the coset format $g\mathbb{G}[\alpha]$. 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 $\mathbb{G}$, i.e. as generated subgraphs $\alpha[1]$ or $\alpha[g]$ in $\mathbb{C}\mathbb{G}$.

A natural way of putting extra constraints on generated subgraphs with reasonable closure properties in terms of generator sets $\alpha \subseteq \mathsf{E}$ is the following. Consider a fixed E-graph $\mathbb{I} = (S, (R_e)_{e\in\mathsf{E}})$ on vertex set S. We want to consider $\mathbb{I}$ as a template for systematic restrictions on patterns of generator sequences, and correspondingly regard $\mathbb{I}$ as a *constraint graph*. With $\mathbb{I}$ we associate the set of all (reduced) words over the alphabet E that label walks in $\mathbb{I}$. A *walk* in $\mathbb{I}$, from a source vertex s to a target vertex t is a finite sequence $s = s_0, e_1, v_1, e_2, \ldots, e_n, s_n = t$ with $e_i \in \alpha$ and $(s_{i-1}, s_i) \in R_{e_i}$.

Definition 3.5. [reduced words over $\mathbb{I}$]
For a constraint graph $\mathbb{I}$ and any subset $\alpha \subseteq \mathsf{E}$ define these sets of reduced words that label walks in $\mathbb{I}$:

$$\begin{array}{rcl} \mathsf{E}^*[\mathbb{I}] & := & \{w \in \mathsf{E}^* \colon\ w \text{ labelling a walk in } \mathbb{I}\ \}, \\ \alpha^*[\mathbb{I}] & := & \mathsf{E}^*[\mathbb{I}] \cap \alpha^*. \end{array}$$

For a finer distinction, sets of reduced words corresponding to walks from a

specified source vertex s, and possibly target vertex t, are similarly defined as

$$\begin{array}{rcl} \mathsf{E}^*[\mathbb{I},s] & := & \{w \in \mathsf{E}^* \colon\ w \text{ labelling a walk from } s \text{ in } \mathbb{I}\ \}, \\ \alpha^*[\mathbb{I},s] & := & \mathsf{E}^*[\mathbb{I},s] \cap \alpha^*, \\ \mathsf{E}^*[\mathbb{I},s,t] & := & \{w \in \mathsf{E}^* \colon\ w \text{ labelling a walk from } s \text{ to } t \text{ in } \mathbb{I}\ \}, \\ \alpha^*[\mathbb{I},s,t] & := & \mathsf{E}^*[\mathbb{I},s,t] \cap \alpha^*. \end{array}$$

Note that

$$\mathsf{E}^*[\mathbb{I}] = \bigcup_{s \in S} \mathsf{E}^*[\mathbb{I},s] = \bigcup_{s,t \in S} \mathsf{E}^*[\mathbb{I},s,t],$$

and that some of the sets $\mathsf{E}^*[\mathbb{I},s,t]$ may be empty, since the constraint graph $\mathbb{I}$ may not be connected. Concatenation of generator sequences associated with the $\mathsf{E}^*[\mathbb{I},s,t]$ will underpin a groupoidal composition operation to be investigated in Section 4. At this point, however, the emphasis is on a notion of $\mathbb{I}$-*reachability* in E-graphs $\mathbb{H}$ that requires a reference just to some specific choice of a source vertex $s \in S$ for the vertex of departure in $\mathbb{H}$: label sequences from $\mathsf{E}^*[\mathbb{I},s]$ single out precisely walks from vertex v in $\mathbb{H}$ that play out an edge sequence of some walk from s in $\mathbb{I}$: $\mathsf{E}^*[\mathbb{I},s]$- *or* $\alpha^*[\mathbb{I},s]$-*walks* in $\mathbb{H}$.

Specifically, in the E-graph $\mathbb{C}\mathbb{G}$, we write $\mathbb{C}\mathbb{G}[\mathbb{I},s,\alpha] \subseteq \mathbb{C}\mathbb{G}[\alpha] \subseteq \mathbb{C}\mathbb{G}$ (and just $\mathbb{C}\mathbb{G}[\mathbb{I},s] \subseteq \mathbb{C}\mathbb{G}$ if $\alpha = \mathsf{E}$) for correspondingly generated weak subgraphs of $\mathbb{C}\mathbb{G}$ at $1 \in G$ on vertex sets

$$\{g \in G \colon G \text{ reachable from } 1 \in G \text{ on a walk labelled by } w \in \alpha^*[\mathbb{I},s]\ \};$$

into the edge sets we include all edges traversed by $\alpha^*[\mathbb{I},s]$-*walks* from 1 in $\mathbb{C}\mathbb{G}$. These weak subgraphs may thus be seen as the homomorphic images in $\mathbb{G}$ of $\alpha^*[\mathbb{I},s] \subseteq \mathsf{E}^*$ under $[\ ]_{\mathbb{G}}$. We also write $\mathbb{G}[\mathbb{I},s,\alpha] \subseteq G$ and $\mathbb{G}[\mathbb{I},s] \subseteq G$ for these sets of group elements but keep in mind that they do not carry the structure of a subgroup.

In any E-graph $\mathbb{H}$, we may similarly look at $\alpha^*[\mathbb{I},s]$-*walks* in $\mathbb{H}$, walks whose edge labelling is in $\alpha^*[\mathbb{I},s]$: the $\alpha^*[\mathbb{I},s]$-component of a vertex v in $\mathbb{H}$ then consist of all those vertices that are reachable from v in $\mathbb{H}$ on $\alpha^*[\mathbb{I},s]$-walks. We write $\alpha^*[\mathbb{I},s;v]$ for the set of vertices in this component, and

$$\mathbb{H}[\mathbb{I},s,\alpha;v] \subseteq \mathbb{H}[\alpha;v] \subseteq \mathbb{H}$$

for the weak subgraph on this component, again with edge sets generated by the edges traversed in $\alpha^*[\mathbb{I},s]$-walks from v in $\mathbb{H}$. These notions naturally generalise the notions of the α-component $\alpha[v]$ and generated subgraph $\mathbb{H}[\alpha;v]$ from Definition 2.7.

In this sense, the set $\mathbb{G}[\mathbb{I},s,\alpha] \subseteq G$ as defined above is the $\alpha^*[\mathbb{I},s]$-component of $1 \in \mathbb{C}\mathbb{G}$, $\mathbb{C}\mathbb{G}[\mathbb{I},s,\alpha] \subseteq \mathbb{C}\mathbb{G}$ the corresponding weak subgraph. Clearly $\alpha^*[\mathbb{I},s]$-components of any elements $g \in \mathbb{C}\mathbb{G}$ support coset-like copies of $\mathbb{C}\mathbb{G}[\mathbb{I},s,\alpha]$ or $\mathbb{C}\mathbb{G}[\mathbb{I},s]$ at $g \in \mathbb{G}$, denoted $g\mathbb{C}\mathbb{G}[\mathbb{I},s,\alpha]$ or $g\mathbb{C}\mathbb{G}[\mathbb{I},s]$, which are naturally obtained as

$$g\mathbb{C}\mathbb{G}[\mathbb{I},s,\alpha] = \mathbb{C}\mathbb{G}[\mathbb{I},s,\alpha;g] \subseteq g\mathbb{C}\mathbb{G}[\alpha] \subseteq \mathbb{C}\mathbb{G},$$

regarded as α-graphs on the vertex sets $\alpha^*[\mathbb{I}, s; g]$ in $\mathbb{C}\mathbb{G}$, which can also be described as the sets

$$\{g[w]_{\mathbb{G}} \colon w \in \alpha^*[\mathbb{I}, s]\} = \{gh \colon h \in \mathbb{G}[\mathbb{I}, s, \alpha]\} \subseteq G.$$

We call these subsets $\mathbb{I}$-*cosets*, even though their algebraic nature as cosets can only be captured in the groupoidal sense to be discussed in Section 4.

We may now generalise the notions of amalgamation chains and coset cycles to the use of α-graphs of the form $\mathbb{C}\mathbb{G}[\mathbb{I}, s, \alpha]$ instead of just $\mathbb{C}\mathbb{G}[\alpha]$. The aim is to obtain E-groups that avoid coset cycles based on overlapping copies of $\mathbb{C}\mathbb{G}[\mathbb{I}, s, \alpha]$.

Definition 3.6. [α-similarity with $\mathbb{G}[\mathbb{I}, s]$]
Let $\mathbb{G}$ be an E-group, $\alpha \subseteq \mathsf{E}$. An E-graph $\mathbb{H} = (V, (R_e)_{e \in \mathsf{E}})$ is *α-similar to $\mathbb{G}[\mathbb{I}, s]$ at $v \in V$* if the α-generated subgraph $\mathbb{H}[\alpha; v]$ of $\mathbb{H}$ is isomorphic to a weak subgraph of $\mathbb{C}\mathbb{G}[\mathbb{I}, s, \alpha] \subseteq \mathbb{C}\mathbb{G}[\mathbb{I}, s]$, via an isomorphism that associates $v \in V$ with $1 \in \mathbb{C}\mathbb{G}$.

Recall that an E-group $\mathbb{G}$ is compatible with the E-graph $\mathbb{I}$ if $[w]_{\mathbb{I}} = \mathrm{id}_S$ for all $w \in \mathsf{E}^*$ such that $[w]_{\mathbb{G}} = 1$. In this case, $g \in G$ induces a well-defined permutation $g^{\mathbb{I}}$ on S that maps $s \in S$ to $g^{\mathbb{I}}(s) := [w]_{\mathbb{I}}(s)$ for any $w \in \mathsf{E}^*$ such that $g = [w]_{\mathbb{G}}$, so that

$$\begin{array}{rcl} \cdot^{\mathbb{I}} \colon \mathbb{G} & \longrightarrow & \mathrm{Sym}(S) \\ g & \longmapsto & g^{\mathbb{I}} \end{array}$$

is a homomorphism from $\mathbb{G}$ onto a subgroup of the symmetric group $\mathrm{Sym}(S)$. Moreover, for any $g \in \alpha^*[\mathbb{I}, s, 1]$ in $\mathbb{C}\mathbb{G}$,

$$\mathbb{C}\mathbb{G}[\mathbb{I}, s, \alpha; 1] = \mathbb{C}\mathbb{G}[\mathbb{I}, s, \alpha] = g\mathbb{C}\mathbb{G}[\mathbb{I}, g^{\mathbb{I}}(s), \alpha] = \mathbb{C}\mathbb{G}[\mathbb{I}, g^{\mathbb{I}}(s), \alpha; g].$$

Definition 3.7. [amalgamation chain over $\mathbb{I}$]
Let $\mathbb{G}$ be compatible with $\mathbb{I}$ and let $(\mathbb{C}\mathbb{G}[\mathbb{I}, s_i, \alpha_i], g_i)_{1 \leqslant i \leqslant n}$ be a sequence of generated subgraphs in $\mathbb{C}\mathbb{G}[\mathbb{I}]$ with distinguished elements $g_i \in \mathbb{G}[I, s_i, \alpha_i]$ such that $g^{\mathbb{I}}(s_i) = s_{i+1}$. Then *amalgamation chains* $\mathbb{H}_k := \oplus_{i=1}^{k}(\mathbb{C}\mathbb{G}[\mathbb{I}, s_i, \alpha_i], g_i)$ are conditionally defined by induction on $1 \leqslant k \leqslant n$, together with distinguished vertices v_k in $\mathbb{H}_k$ according to:

(i) $\mathbb{H}_1, v_1 := \mathbb{C}\mathbb{G}[\mathbb{I}, s_1, \alpha_1], g_1$ for $k = 1$ (unconditionally);

(ii) for $k < n$, and under the condition that $\alpha_{k+1}[\mathbb{I}, s_{k+1}; v_k] \subseteq \alpha_k[\mathbb{I}, s_{k+1}; v_k]$ in $\mathbb{H}_k$, let

$$\mathbb{H}_{k+1} := (\mathbb{H}_k, v_k) \oplus (\mathbb{C}\mathbb{G}[\mathbb{I}, s_{k+1}, \alpha_{k+1}], 1)$$

and v_{k+1} the vertex corresponding to g_{k+1} in the amalgamated copy of $\mathbb{C}\mathbb{G}[\mathbb{I}, s_{k+1}, \alpha_{k+1}]$.

Definition 3.8. [$\mathbb{I}$-coset cycles and N-acyclicity]
For $n \geqslant 2$, an $\mathbb{I}$-*coset cycle of length n* in an E-group $\mathbb{G}$ that is compatible with $\mathbb{I}$, is a cyclically indexed tuple of pointed $\mathbb{I}$-cosets $(g_i\mathbb{G}[\mathbb{I}, s_i, \alpha_i], g_i)_{i \in \mathbb{Z}_n}$ such that, for all i,

(i) $s_{i+1} = h_i^{\mathbb{I}}(s_i)$, where $h_i := g_i^{-1} g_{i+1}$;
(ii) $g_{i+1} \in g_i \mathbb{G}[\mathbb{I}, s_i, \alpha_i]$;
(iii) $g_i \mathbb{G}[\mathbb{I}, s_i, \alpha_i \cap \alpha_{i-1}] \cap g_{i+1} \mathbb{G}[\mathbb{I}, s_{i+1}, \alpha_i \cap \alpha_{i+1}] = \emptyset$.

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.

As before, the special case of 2-acyclicity over $\mathbb{I}$ can also be regarded as a notion of simple connectivity w.r.t. $\mathbb{I}$-cosets or $\mathbb{I}$-reachability.

Note that condition (i) in the definition makes sure that the action of $\alpha_i^*[\mathbb{I}, s_i]$ carries $1 \in \mathbb{G}[\mathbb{I}, s_i, \alpha_i]$, which is associated with $s_i \in S$ to $h_i := g_i^{-1} g_{i+1} \in \mathbb{G}[\mathbb{I}, s_i, \alpha_i]$, which is associated with $s_{i+1} \in S$ (in $\mathbb{G}[\mathbb{I}, s_i, \alpha_i]$ and then, for the next step, in $\mathbb{G}[\mathbb{I}, s_{i+1}, \alpha_{i+1}]$).

A slightly different, yet instructive perspective on $\mathbb{I}$-cosets and related generated subsets of $\mathbb{C}\mathbb{G}$ is in terms of $\mathbb{I}$-*reachability* in an S-annotated version of $\mathbb{C}\mathbb{G}$. If $\mathbb{G}$ is compatible with $\mathbb{I}$, then any choice of an element $s \in S$ to annotate some specific group element g extends to a unique S-colouring of all of $\mathbb{G}$ that associates the colour $s' := h^{\mathbb{I}}(s)$ with the group element $g' = gh$. The side conditions on $\mathbb{I}$-coset cycles w.r.t. the $s_i \in S$ involved (item (i) in the definition above) imply that all the $\mathbb{I}$-cosets involved adhere to the same S-colouring; they refer to $\mathbb{I}$-reachability (within α_i-generated cosets of $\mathbb{G}$) with reference to that common annotation of group elements with elements of S.

Instead of considering individual S-annotations of an $\mathbb{I}$-compatible E-group $\mathbb{G}$, we may also consider all these admissible annotations in parallel. This leads to the following substructure of the cartesian product structure $\mathbb{I} \times \mathbb{C}\mathbb{G}$. The latter is the usual direct product of the two E-graphs $\mathbb{I}$ and $\mathbb{C}\mathbb{G}$ as relational structures: the vertex set is $S \times G$ and $((s, g), (s', g')) \in R_e$ if $(s, s') \in R_e^{\mathbb{I}}$ and $(g, g') \in R_e^{\mathbb{C}\mathbb{G}}$.

Definition 3.9. [$\mathbb{I}$-products]
For an E-group $\mathbb{G}$ that is compatible with the constraint graph $\mathbb{I}$ we define the $\mathbb{I}$-*product* $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ to be the weak substructure of $\mathbb{I} \times \mathbb{C}\mathbb{G}$ that is made up of the $\mathsf{E}^*[\mathbb{I}, s]$-components of the elements $(s, 1) \in S \times G$ in $\mathbb{I} \times \mathbb{C}\mathbb{G}$.

In other words, $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ is the closure of the subset $S \times \{1\}$ under $\mathbb{I}$-reachability, where the tag $s \in S$ in the first component is the annotation for $1 \in G$, which implies that

$$\mathbb{I} \otimes \mathbb{C}\mathbb{G} \;\simeq\; \dot{\bigcup}_{s \in S} \mathbb{C}\mathbb{G}_s \quad \text{for } \mathbb{C}\mathbb{G}_s := \mathbb{C}\mathbb{G}[\mathbb{I}, s].$$

The first of the following remarks casts N-acyclicity of $\mathbb{G}$ over $\mathbb{I}$ in terms of $\mathbb{I}$-reachability in $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ of its component structures $\mathbb{C}\mathbb{G}_s$; the sattement itself is obvious from Definitions 3.8 and 3.9. The second remark concerns the obvious symmetries among these component structures.

Remark 3.10. *Let the E-group $\mathbb{G}$ be compatible with the constraint graph $\mathbb{I} = (S, \mathsf{E})$ and N-acyclic over $\mathbb{I}$. Then $\mathbb{H} := \mathbb{I} \otimes \mathbb{C}\mathbb{G}$ has no cycles of length $n \leqslant N$*

formed by non-trivially overlapping α_i-components of vertices $v_i = (s_i, g_i)$ of $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ of the form

$$(\mathbb{H}[\alpha_i; v_i])_{i \in \mathbb{Z}_n}$$

where the subsets $\alpha_i \subseteq \mathsf{E}$ and vertices $v_i = (s_i, g_i)$ are such that, for $i \in \mathbb{Z}_n$,

(i) $v_{i+1} \in \mathbb{H}[\alpha_i; v_i] \cap \mathbb{H}[\alpha_{i+1}; v_{i+1}]$, *and*

(ii) $\mathbb{H}[\alpha_i \cap \alpha_{i-1}; v_i] \cap \mathbb{H}[\alpha_i \cap \alpha_{i+1}; v_{i+1}] = \emptyset$.

Remark 3.11. *Let the E-group $\mathbb{G}$ be compatible with the constraint graph $\mathbb{I} = (S, \mathsf{E})$. Then the $\mathbb{I}$-product $S \otimes_{\mathbb{I}} \mathbb{C}\mathbb{G}$ is isomorphic to the disjoint union of S-annotated E-graphs $\mathbb{C}\mathbb{G}_s$, where $\mathbb{C}\mathbb{G}_s$ is $\mathbb{C}\mathbb{G}[\mathbb{I}, s]$ with the unique $\mathbb{I}$-induced vertex colouring that associates colour $s \in S$ with $1 \in G$. Note that the component E-graphs $\mathbb{C}\mathbb{G}_s$ and $\mathbb{C}\mathbb{G}_t$ are isomorphic whenever s and t are in the same connected component of $\mathbb{I}$. If the constraint graph $\mathbb{I}$ is connected, then all components $\mathbb{C}\mathbb{G}_s$ share the same isomorphism type.*

For the isomorphism claim assume that t is reachable from s in $\mathbb{I}$. This implies that $g^{\mathbb{I}}(s) = t$ for some suitable g, so that g is coloured by t in $\mathbb{C}\mathbb{G}_s$. It follows that left multiplication with g^{-1}, which is an automorphism of the E-graph $\mathbb{C}\mathbb{G}$ mapping g to 1, transforms $\mathbb{C}\mathbb{G}_s$ into $\mathbb{C}\mathbb{G}_t$.

The following is a strict analogue to Lemma 2.15 for amalgamation chains over $\mathbb{I}$.

Lemma 3.12. *Provided the constituents $\mathbb{C}\mathbb{G}[\mathbb{I}, s_i, \alpha_i]$ are 2-acyclic over $\mathbb{I}$, all β-components in an amalgamation chain $\mathbb{H} = \bigoplus_{i=1}^{n} (\mathbb{C}\mathbb{G}[\mathbb{I}, s_i, \alpha_i], g_i)$ arise as amalgamation chains of lengths up to n based on constituents $\mathbb{C}\mathbb{G}[\mathbb{I}, s_i, \alpha_i \cap \beta]$.*

Proof. The proof is strictly analogous to that of Lemma 2.15, with the only difference that all the generated α_i-, β- and $(\alpha_i \cap \beta)$-components in question are now components over $\mathbb{I}$. □

Proposition 3.13. *For every finite set E, finite E-graph $\mathbb{I}$ and $N \geqslant 2$, there are finite E-groups $\mathbb{G}$ that are compatible with any amalgamation chains of lengths up to N generated by any subsets $\mathbb{C}\mathbb{G}[\mathbb{I}, s_i, \alpha_i]$. In particular, such $\mathbb{G}$ is guaranteed to be N-acyclic over $\mathbb{I}$. Moreover, the E-groups obtained in the proposed construction are fully symmetric w.r.t. all those permutations of the set E that are also symmetries of $\mathbb{I}$. One can also obtain finite E-groups with the above property that are compatible with another given finite E-graph $\mathbb{H}$, and symmetric w.r.t. all those permutations of the set E that are symmetries of both the E-graphs $\mathbb{I}$ and $\mathbb{H}$.*

Proof. Analogous to Proposition 3.4, by induction on $n \leqslant |\mathsf{E}|$, construct $\mathbb{G}_n$ such that all generated subgroups $\mathbb{G}_n[\alpha]$ for $|\alpha| \leqslant n$ are compatible with $\mathbb{I}$-amalgamation chains of lengths up to N, and $\mathbb{G}_n[\alpha] \simeq \mathbb{G}_m[\alpha]$ for $n, m \geqslant |\alpha|$. Again $\mathbb{G}_{n+1}$ can be obtained as $\mathrm{sym}(\mathbb{H}_n)$, i.e. generated by permutation group action on an E-graph $\mathbb{H}_n$ that contains as components all $\mathbb{I}$-amalgamation chains of length up to N of $\mathbb{I}$-cosets generated by generator subsets of sizes up to n. □

Remark 3.14. *Levels of acyclicity over $\mathbb{I}$ can be cast so that they in particular imply N-coset-acyclicity in the sense of Definition 3.2: for this simply augment the given $\mathbb{I}$ by a separate 1-vertex component whose only vertex has R_e-loops for all $e \in \mathsf{E}$. In this sense, Proposition 3.13 serves as a generalisation that comprises Proposition 3.4 as a special case. Moreover, and in the same vein, any finite number of finite constraint graphs $\mathbb{I}$ can in effect be replaced by their disjoint union in order to cover, in a single constraint graph $\mathbb{I}$, all relevant cyclic patterns.*

4 Groupoidal variants

In terms of the combinatorial action of the generators $e \in \mathsf{E}$ on an E-graph $\mathbb{H}$, and by extension of the monoid structure E^* on $\mathbb{H}$, the involutive nature of $\pi_e \in \mathrm{Sym}(V)$ is closely tied to the undirected nature of e-edges in E-graphs. We want to overcome this constraint by allowing for directed e-edges. At the same time we may want to relax the strictly prescribed uniformity between vertices, as we have already achieved in the context of involutive generators with constraint graphs $\mathbb{I}$ in Section 3.3. So we want to allow for vertices of different sorts with directed transitions via e-edges between vertices of specific sorts. Indeed, some applications of related notions of acyclicity in graph and hypergraph structures inspired by the idea of Cayley graphs in [11, 13] can be naturally cast in terms of such multi-sorted multi-graph structures and related groupoids. The novelty of their treatment here is that we can directly reduce the construction of groupoids with the desired coset acyclicity properties to the simpler constructions of groups from the previous section.

4.1 Directed multi-graphs as constraint patterns

In the following we want to consider groupoid structures with a specified pattern of sorts (types of elements, objects) and generators (for the groupoidal operation, morphisms). Groupoids in our sense can also be associated with inverse semigroups of correspondingly restricted pattern. 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* – indeed the corresponding structures generalise the constraint graphs of Section 3.3 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 $e \in \mathsf{E}$ are directed, with an explicit operation of edge reversal. as associated with inverse semigroups,

Definition 4.1. [constraint pattern $\mathbb{I}$]
A *constraint pattern* is a multi-graph $\mathbb{I} = (S, \mathsf{E}, \iota_1, \iota_2, \cdot^{-1})$, which we formalise as a two-sorted structure with a set S of vertices and a set E of edges as sorts, linked by maps $\iota_i \colon \mathsf{E} \to S$ that associate a source and target vertex with every edge $e \in \mathsf{E}$, 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 $\iota_1(e^{-1}) = \iota_2(e)$.

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

In order to extend the notion of $\mathbb{I}$-reachability (based on an undirected constraint graph $\mathbb{I}$) to a similar concept of $\mathbb{I}$-reachability w.r.t. a constraint pattern $\mathbb{I}$ that is a directed multi-graph, we consider words that label directed walks in $\mathbb{I}$. For the following compare Definition 3.5 and related notions in Section 3.3. A *reduced word* over E now is a word in which no $e \in \mathsf{E}$ is directly followed or preceded by its inverse e^{-1}. Similar to the terminology in Section 3.3, we write $\mathsf{E}^*[\mathbb{I}]$ for the set of all reduced words over E that label walks in $\mathbb{I}$, and naturally extend the ι-maps to all of $\mathsf{E}^*[\mathbb{I}]$ as follows. Since a walk from s to t in $\mathbb{I}$ is a sequence $s = s_0, e_1, s_1, \ldots, e_n, s_n = t$ such that $\iota_1(e_i) = s_{i-1}$ and $\iota_2(e_i) = s_i$ for $1 \leqslant i \leqslant n$, this walk is fully determined by the sequence of edges and can be identified with the E-word $w = e_1 \ldots e_n$. So we think of $\mathsf{E}^*[\mathbb{I}]$ as the set of all E-words $w = e_1 \ldots e_n$ with $\iota_2(e_i) = \iota_1(e_{i+1})$ for $1 \leqslant i < n$, and put $\iota_1(w) := \iota_1(e_1)$ and $\iota_2(w) := \iota_2(e_n)$, so that w labels a walk in $\mathbb{I}$ from the source vertex $\iota_1(w)$ to the target vertex $\iota_2(w)$. Correspondingly we now define

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

so that $\mathsf{E}[\mathbb{I}] = \dot{\bigcup}_{s,t \in S} \mathsf{E}^*[\mathbb{I}, s, t]$.

Concatenation between words or walks w_1 and w_2 is defined as a walk $w_1 w_2 \in \mathsf{E}^*[\mathbb{I}, \iota_1(w_1), \iota_2(w_2)]$ whenever their ι-values match in the sense that $\iota_2(w_1) = \iota_1(w_2)$. All these notions are similarly available in restriction to any subset $\alpha \subseteq \mathsf{E}$ that is closed under edge reversal; for instance $\alpha^*[\mathbb{I}, s, t] \subseteq \alpha^*[\mathbb{I}]$ are defined as corresponding sets of reduced α-words in relation to the pattern $\mathbb{I}$.

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

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

Definition 4.2. [$\mathbb{I}$-groupoid]
An $\mathbb{I}$-*groupoid* based on the constraint pattern $\mathbb{I} = (S, \mathsf{E})$ is a groupoid structure of the form $\mathbb{G} = (G, (G_{s,t})_{s,t \in S}, \cdot, (1_s)_{s \in S}, (g_e)_{e \in \mathsf{E}})$ where

(i) the family $(G_{s,t})_{s,t \in S}$ partitions the universe G of groupoid elements;[2]

(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 all $s \in S$;

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

(iv) G is generated by the family of pairwise distinct elements $g_e \in G_{\iota_1(e),\iota_2(e)}$ for $e \in \mathsf{E}$, 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 corresponding 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 4.3. [Cayley graph of an $\mathbb{I}$-groupoid]
The *Cayley graph* of an $\mathbb{I}$-groupoid $\mathbb{G} = (G, (G_{s,t})_{s,t\in S}, \cdot, (1_s)_{s\in S}, (g_e)_{e\in\mathsf{E}})$ is the directed edge-coloured graph $\mathbb{C}\mathbb{G} := \mathrm{Cayley}(\mathbb{G}) = (G, (R_e)_{e\in\mathsf{E}})$ with vertex set G and edge sets of colour $e \in \mathsf{E}$ according to

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

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

$$[w]_{\mathbb{G}} = g_{e_1} \cdots g_{e_n} \in G_{s,t}.$$

It is clear from the definition that $g = [w]_{\mathbb{G}} \in G_{s,t}$ precisely for $s = \iota_1(w)$ and $t = \iota_2(w)$. We may therefore consistently define ι-maps directly for the elements of an $\mathbb{I}$-groupoid $\mathbb{G}$ according to

$$\begin{array}{rcl} \iota_i \,:\, G & \longrightarrow & S \\ g \in G_{s,t} & \longmapsto & \begin{cases} \iota_1(g) := s, \\ \iota_2(g) := t. \end{cases} \end{array}$$

Similarly writing $\alpha^*[\mathbb{I}, s, t]$ for the set of (reduced) words over a subset $\alpha \subseteq \mathsf{E}$ that is closed under edge reversal,

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

we can look at generated substructures in groupoids or their Cayley graphs. In particular, generated sub-groupoids $\mathbb{G}[\alpha]$, for subsets $\alpha \subseteq \mathsf{E}$ that are closed under edge reversal, and corresponding groupoidal cosets at $g \in G$, are defined in the obvious manner as

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

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

As with Cayley graphs for E-groups, the Cayley graphs of $\mathbb{I}$-groupoids are more homogeneous than the underlying groupoid, simply because groupoidal composition is only encoded in terms of right multiplication with individual generators; in particular, the neutral elements 1_s can in general not be identified in $\mathbb{CG}$. What is still recognisable in $\mathbb{CG}$ for an $\mathbb{I}$-groupoid $\mathbb{G}$, is membership in the sets

$$G[*,t] := \bigcup_{s\in S} G_{s,t} = \{g \in G \colon \iota_2(g) = t\}$$

for each $t \in S$ that is not isolated in $\mathbb{I}$: this is the set of vertices with an outgoing R_e-edge for any e with $\iota_1(e) = t$. So Cayley graphs of $\mathbb{I}$-groupoids are not as homogeneous as Cayley graphs of groups, simply because groupoidal composition is not total but requires matching sorts.

The algebraic structure of the $\mathbb{I}$-groupoid $\mathbb{G}$ is, however, still fully determined and can be recovered from its Cayley graph $\mathbb{CG}$ in the corresponding action of partial permutations. 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, we here realise the groupoid as a subgroupoid of the set of all bijections between the relevant sets $G[*,t]$.

Observation 4.4. *The $\mathbb{I}$-groupoid $\mathbb{G}$ is isomorphic to the $\mathbb{I}$-groupoid generated by the following partial bijections π_e for $e \in \mathsf{E}[t,t']$:*

$$\begin{array}{rcl} \pi_e \colon G[*,t] & \longrightarrow & G[*,t'] \\ g & \longmapsto & g \cdot g_e, \end{array}$$

where $g \cdot g_e$ is identified as the unique vertex g' of $\mathbb{CG}$ for which $(g,g') \in R_e$. Here 1_t is the partial bijection $\mathrm{id}_s \colon G[*,t] \to G[*,t]$.

For the following compare Definitions 2.1 and 2.8.

Definition 4.5. [$\mathbb{I}$-graphs and compatibility]
For a constraint pattern $\mathbb{I} = (S, \mathsf{E})$, an $\mathbb{I}$-*graph* is a vertex- and edge-coloured directed graph $\mathbb{H} = (V, (V_s)_{s\in S}, (R_e)_{e\in\mathsf{E}})$, 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 $e \in \mathsf{E}$ such that $R_{e^{-1}} = R_e^{-1}$; it is *complete* if each R_e is a complete matching between $V_{\iota_1(e)}$ and $V_{\iota_2(e)}$, i.e. is the graph of a bijection $\pi_e \colon V_{\iota_1(e)} \to V_{\iota_2(e)}$.

In a complete $\mathbb{I}$-graph $\mathbb{H}$, the composition of the π_e along a walk $w \in \mathsf{E}^*[\mathbb{I}, s, t]$ induces a bijection $\pi_w \colon V_s \to V_t$, which we also denote as $[w]_{\mathbb{H}}$. An $\mathbb{I}$-groupoid $\mathbb{G}$ is *compatible* with the complete $\mathbb{I}$-graph $\mathbb{H}$ if for all $w \in \mathsf{E}^*[\mathbb{I}, s, s]$

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

The following are straightforward illustrations of these concepts. We note that the Cayley graph $\mathbb{CG}$ of an $\mathbb{I}$-groupoid $\mathbb{G}$ can be cast as a complete $\mathbb{I}$-graph if we colour each vertex $g \in G$ by $\iota_2(g) \in S$.

Observation 4.6. *Any $\mathbb{I}$-groupoid $\mathbb{G}$ is compatible with its Cayley graph if we view the latter as a complete $\mathbb{I}$-graph in the natural manner. Another $\mathbb{I}$-groupoid $\hat{\mathbb{G}}$ is compatible with the Cayley graph $\mathbb{C}\mathbb{G}$ of $\mathbb{G}$ if, and only if, there is a homomorphism $h\colon \hat{\mathbb{G}} \to \mathbb{G}$, if, and only if, $h\colon [w]_{\hat{\mathbb{G}}} \mapsto [w]_{\mathbb{G}}$ for $w \in \mathsf{E}^*[\mathbb{I}]$ is well-defined as a map from $\hat{\mathbb{G}}$ to $\mathbb{G}$.*

4.2 Coset acyclicity for groupoids

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

Definition 4.7. [coset cycles]
For $n \geqslant 2$, a *coset cycle of length n* in an $\mathbb{I}$-groupoid $\mathbb{G}$ is a cyclically indexed tuple of pointed cosets $(g_i\mathbb{G}[\alpha_i], g_i)_{i\in\mathbb{Z}_n}$ such that, for all i,

(i) $g_{i+1} \in g_i\mathbb{G}[\alpha_i]$;

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

Definition 4.8. [N-acyclicity]
For $N \geqslant 2$, an $\mathbb{I}$-groupoid $\mathbb{G}$ is *N-coset-acyclic* (*N-acyclic* for short) if it admits no coset cycles of lengths up to N.

4.3 From groups to groupoids

With a constraint pattern $\mathbb{I} = (S, \mathsf{E})$ for $\mathbb{I}$-groupoids $\mathbb{G}$ we associate a richer set $\hat{\mathsf{E}}$ of involutive generators for $\hat{\mathsf{E}}$-groups $\hat{\mathbb{G}}$ so that interesting $\mathbb{I}$-groupoids $\mathbb{G}$ can be identified within suitable $\hat{\mathsf{E}}$-groups $\hat{\mathbb{G}}$; more specifically, we aim for a low-level interpretation of Cayley graphs of $\mathbb{I}$-groupoids $\mathbb{C}\mathbb{G}$ within products of the Cayley graph $\mathbb{C}\hat{\mathbb{G}}$ of suitable $\hat{\mathsf{E}}$-groups $\hat{\mathbb{G}}$ with a suitable constraint graph $\hat{\mathbb{I}}$.

Firstly, in order to interpret the directed multi-graph structure of

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

in an $\hat{\mathsf{E}}$-graph structure $\hat{\mathbb{I}} = (\hat{S}, (R_{\hat{e}})_{\hat{e}\in\hat{\mathsf{E}}})$, we associate 3 new generators with every generator $e \in \mathsf{E}$, and insert two new vertices into $\hat{S} \supseteq S$ so as to represent directed e-(multi-)edges as paths of length 3 in $\hat{\mathbb{I}}$, as follows.

A directed edge $e \in \mathsf{E}[s, s']$ in $\mathbb{I}$ and its inverse $e' := e^{-1} \in \mathsf{E}[s', s]$ are to be replaced by a succession of 3 undirected edges with labels $\{e\}$, $\{e, e'\}$ and $\{e'\}$ with 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 $e \in \mathsf{E}[s, s]$ at s and its inverse $e' := e^{-1}$, correspondingly get replaced by a cycle of 3 undirected edges with labels $\{e\}$, $\{e, e'\}$ and $\{e'\}$.

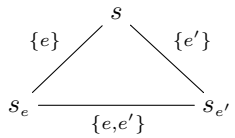

Note that these replacements are inherently symmetric w.r.t. edge reversal in the sense that the replacements really concern the edge pair $\{e, e^{-1}\}$. The direction of e is encoded in the directed nature of the walk

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

whose reversal exactly is the corresponding walk for e^{-1}.

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

Let $\hat{\mathsf{E}}$ be the set of these new edge labels, $\hat{S}$ the set S of vertices of $\mathbb{I}$ together with the 2 newly introduced intermediate vertices for each $e \in \mathsf{E}$, and $\hat{\mathbb{I}}$ the $\hat{\mathsf{E}}$-graph on vertex set $\hat{S}$ resulting from the replacement scheme above. Note that in $\hat{\mathbb{I}}$, just as in $\mathbb{I}$, every edge label occurs exactly once, and $\hat{\mathbb{I}}$ is a simple undirected graph without loops. Moreover, for $s, t \in S \subseteq \hat{S}$ there is a one-to-one correspondence between reduced words in

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

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

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

is a bijection. For this observation it is essential that reduced words in $\hat{\mathsf{E}}^*[\hat{\mathbb{I}}]$ 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 $e \in \mathsf{E}$ with the triplet for e^{-1} would not be a reduced $\hat{\mathsf{E}}$-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 $\hat{\mathsf{E}}$-word can leave this $\{\{e\}, \{e, e^{-1}\}, \{e^{-1}\}\}$-component of $\hat{\mathbb{I}}$ is via $\iota_1(e)$ or $\iota_2(e)$, which are both in S.

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

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

Interestingly, we may now extract an $\mathbb{I}$-groupoid $\mathbb{G}$ from any $\hat{\mathsf{E}}$-group $\hat{\mathbb{G}}$ that is compatible with the constraint graph $\hat{\mathbb{I}}$. For that, recall from Definition 3.9 the direct product $\hat{\mathbb{I}} \times \mathbb{C}\hat{\mathbb{G}}$ with the $\hat{\mathbb{I}}$-product $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ as a weak substructure, which represents components w.r.t. $\hat{\mathbb{I}}$-reachability (cf. Remark 3.11). In fact one may think of the Cayley graph of the target groupoid $\mathbb{G} := \hat{\mathbb{G}}[\mathbb{I}]$ as interpreted within

$\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$. The rough idea is to single out those vertices of $\mathbb{C}\hat{\mathbb{G}}[\hat{\mathbb{I}}, s]$-components in $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ that are annotated with values $t \in S$ and to replace $\{e\}\{e, e^{-1}\}\{e^{-1}\}$-paths of length 3 by directed E-edges, thus reversing the translation from E to $\hat{\mathsf{E}}$.

So now define $\mathbb{G}$ in terms of its generators $e \in \mathsf{E}$, which are interpreted as partial bijections on the vertex set of $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$. That vertex set consists of all pairs $(\hat{s}, [\hat{w}]_{\hat{\mathbb{G}}}) \in \hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ with $\hat{w} \in \hat{\mathsf{E}}^*[\mathbb{I}, \hat{s}]$ In this vertex set, we isolate elements from the components $\mathbb{C}\hat{\mathbb{G}}_s = \{(s, [\hat{w}]_{\hat{\mathbb{G}}}) \in \hat{\mathbb{I}} \otimes \hat{G} \colon \hat{w} \in \hat{\mathsf{E}}^*[\hat{\mathbb{I}}, s]\}$ for $s \in S$ whose annotation in terms of $\hat{\mathbb{I}}$-reachability also falls into $S \subseteq \hat{S}$:

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

where the second equality appeals to the identification of reduced words in $\hat{\mathsf{E}}^*[\hat{\mathbb{I}}, t, s]$ and $\mathsf{E}^*[\mathbb{I}, t, s]$ for $s, t \in S \subseteq \hat{S}$. We write $G_{*,s}$ for the union

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

With $e \in \mathsf{E}[t, t']$ we associate the following partial bijection on the vertex set of $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$, with domain and image as indicated:

$$\begin{array}{rcl} g_e \colon G_{*,t} & \longrightarrow & G_{*,t'} \\ (s, \hat{g}) & \longmapsto & (s, \hat{g} \cdot \{e\} \cdot \{e, e^{-1}\} \cdot \{e^{-1}\}) \text{ where } \hat{g}^{\hat{\mathbb{I}}}(s) = t, \text{ i.e.} \\ (s, [\hat{w}]_{\hat{\mathbb{G}}}) & \longmapsto & (s, [\hat{w}\hat{e}]_{\hat{\mathbb{G}}}) \text{ where } w \in \mathsf{E}^*[\mathbb{I}, s, t]. \end{array}$$

By compatibility of $\hat{\mathbb{G}}$ with $\hat{\mathbb{I}}$, the sets $G_{s,t}$ are disjoint subsets of the vertex set of $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ and thus partition

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

into subsets (not all necessarily non-empty unless $\mathbb{I}$ is connected). Over these sets, concatenation (and reduction) of corresponding words or walks in $\hat{\mathbb{I}}$ induces a well-defined groupoid operation according to

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

where the concatenation relies on the condition that $\iota_2(\hat{w}_1) = t = \iota_1(\hat{w}_2)$, which for the words in question is equivalent with $\iota_2(w_1) = t = \iota_1(w_2)$. The neutral element in $G_{s,s}$ is $1_s := (s, [\lambda]_{\hat{\mathbb{G}}})$. With these stipulations,

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

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

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

Moreover, the induced homomorphism

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

is the valuation map for generator sequences w in the groupoid $\mathbb{G}$.

Observation 4.9. *Let $\hat{\mathbb{G}}$ be a $\hat{\mathsf{E}}$-group that is compatible with $\hat{\mathbb{I}}$. Then the Cayley graph of the $\mathbb{I}$-groupoid $\mathbb{G} = \hat{\mathbb{G}}(\mathbb{I})$ as just constructed from $\mathbb{C}\hat{\mathbb{G}}$ is isomorphic, as an E-graph, to the weak subgraph of $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$ whose vertices are the elements of G, and with e-edges represented by $\hat{e} = \{e\}\{e, e^{-1}\}\{e^{-1}\}$-paths of length 3 between vertices from G in $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$.*

If the $\hat{\mathsf{E}}$-group $\hat{\mathbb{G}}$ is symmetric w.r.t. permutations of the set $\hat{\mathsf{E}}$ that are symmetries of $\hat{\mathbb{I}}$ (i.e., that arise from symmetries of the constraint pattern $\mathbb{I}$), then the $\mathbb{I}$-groupoid $\mathbb{G}$ extracted from $\hat{\mathbb{G}}$ shares all these symmetries.

The following is the main technical result of this section, since it reduces the construction of N-coset acyclic groupoids to the simpler construction of Cayley groups with involutive generators that are N-acyclic over some constraint graph. As such, this reduction can in particular replace the much more intricate, stand-alone construction of N-acyclic groupoids from [11].

Proposition 4.10. *Let $\hat{\mathbb{G}}$ be an $\hat{\mathsf{E}}$-groupoid that is compatible with $\hat{\mathbb{I}}$ and N-acyclic over the constraint graph $\hat{\mathbb{I}}$. Then the induced $\mathbb{I}$-groupoid $\mathbb{G} = \hat{\mathbb{G}}(\mathbb{I})$ is N-coset-acyclic. If $\hat{\mathbb{G}}$ is also compatible with the $\hat{\mathsf{E}}$-translation of a complete $\mathbb{I}$-graph $\mathbb{H}$, then $\mathbb{G}$ is compatible with $\mathbb{H}$.*

The main claim, concerning N-acyclicity, follows directly from the following compatibility of the corresponding notions of cycles with the interpretation of $\mathbb{G}$ in $\hat{\mathbb{G}}$, as expressed in the following lemma. The argument towards compatibility with a given $\mathbb{H}$ straightforward.

Lemma 4.11. *There is a natural translation of coset cycles in the groupoid $\mathbb{G} = \hat{\mathbb{G}}(\mathbb{I})$ based on the map $\hat{}$ for generator sets, which translates coset cycles in the groupoid $\mathbb{G}$ to cycles of the same length over $\hat{\mathbb{I}}$ in the group $\hat{\mathbb{G}}$.*

Proof. Let

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

be a coset cycle in the groupoid $\mathbb{G}$, according to Definition 4.7. For $s := \iota_1(g_0)$, all groupoid elements in this cycle are in $G_{s,*} = \bigcup_{t\in S} G_{s,t}$, so that, in terms of the interpretation of $\mathbb{C}\mathbb{G}$ in $\hat{\mathbb{I}} \otimes \mathbb{C}\hat{\mathbb{G}}$, the entire coset cycle is represented in the component $\mathbb{C}\hat{\mathbb{G}}_s$. For $i \in \mathbb{Z}_n$, let $(s, \hat{g}_i) \in \mathbb{C}\hat{\mathbb{G}}_s$ be the vertex representing g_i. As the natural translation of the cycle $(*)$ into $\hat{\mathbb{G}}$, consider

$$(**) \qquad (\hat{g}_i\hat{\mathbb{G}}[\hat{\alpha}_i, s_i], \hat{g}_i)_{i\in\mathbb{Z}_n} \text{ where } s_i = \iota_2(g_{i-1}).$$

This translation in effect replaces the embeddings of the subsets $g_i\mathbb{G}[\alpha]$ in $\mathbb{C}\mathbb{G}_s$ by their closures $\hat{\mathbb{G}}[\hat{\mathbb{I}}, \hat{\alpha}_i, s_i]$ w.r.t. $\hat{\mathbb{I}}$-reachability inside their $\hat{\alpha}_i$-coset.

The closure is obtained as the union of all $\hat{\mathbb{G}}[\hat{\mathbb{I}}, \{\{e\}, \{e, e^{-1}\}, \{e^{-1}\}\}, \iota_2(g_i)]$-components in $\mathbb{C}\hat{\mathbb{G}}_s$ that contain at least one element of $g_i\mathbb{G}[\alpha]$. 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 3.8. In particular, the first condition, that $s_{i+1} = h_i^{\hat{\mathbb{I}}}(s_i)$, where $h_i = g_1^{-1}g_{i+1}$, follows directly from the corresponding condition on $(*)$, that $g_{i+1} \in g_i\mathbb{G}[\alpha_i]$ together with compatibility of $\hat{\mathbb{G}}$ with $\hat{\mathbb{I}}$, which guarantees that $h_i^{\hat{\mathbb{I}}}$ is well-defined over $\hat{S}$ and maps s_i to s_{i+1}: $g_{i+1} = g_i h_i$ in $\mathbb{G}$ implies that $h_i \in G_{s_i,s_{i+1}}$ for $s_i = \iota_2(g_i)$ and $s_{i+1} = \iota_2(g_{i+1})$.

The crucial element in the relevant definitions is the intersection condition, condition (iii) in Definition 3.8 for $(**)$, which follows from condition (ii) in Definition 4.7 for $(*)$.

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

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

By the intersection condition for $(*)$, $\hat{g} \notin$ so that $\iota_2(\hat{g}) \in \hat{S} \setminus S$, i.e. $\iota_2(\hat{g}) \in \{s_e, s_{e^{-1}}\}$ for some $e \in \mathsf{E}$. But in $\hat{\mathbb{G}}[\hat{\mathbb{I}}, \{\{e\}, \{e, e^{-1}\}, \{e^{-1}\}\}]$-components of elements of G, any vertex with ι_2-value outside S is isolated in $\mathbb{C}\hat{\mathbb{G}}_s$ from all vertices in G by $\{e\}$- and $\{e^{-1}\}$-edges (just as vertices in $\hat{S} \setminus S$ are isolated from S in $\hat{\mathbb{I}}$). So $(\ddagger)$ implies that $e, e^{-1} \in \alpha_{i-1} \cap \alpha_i \cap \alpha_{i+1}$, which would imply that there also is an e-link between the elements of that component that are in G, and hence in $g_i\mathbb{G}[\alpha_i \cap \alpha_{i-1}]$ and in $g_{i+1}\mathbb{G}[\alpha_i \cap \alpha_{i+1}]$. This would violate the intersection condition for $(*)$. □

Corollary 4.12. *For any constraint pattern $\mathbb{I} = (S, \mathsf{E})$, any complete $\mathbb{I}$-graph $\mathbb{H}$, and $N \geqslant 2$ there are finite N-acyclic $\mathbb{I}$-groupoids $\mathbb{G}$ that are 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 permutation of E that is also a symmetry of $\mathbb{I}$ and $\mathbb{H}$ gives rise to an automorphism of the $\mathbb{I}$-groupoid $\mathbb{G}$.*

5 Conclusion

The generic constructions of the preceding chapters show the remarkable 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 undirected, 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 contributed 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-coset-acyclic groups with involutive generators from [11] to yield a considerable simplification of

corresponding constructions for groupoids from [12]. Due to the symmetry preserving, generic character, the new presentation also supports the use of these groupoids in [13], where symmetry considerations are of the essence towards lifting local to global symmetries in finite structures.

To conclude our new self-contained treatment of these group(oid) constructions here, let us briefly look at the most salient application for finite groups and groupoids of graded coset-acyclicity: the construction of finite coverings of graphs and hypergraphs that unravel short cycles:

(1) Natural, unbranched finite coverings of graphs by graphs with interesting acyclicity properties can be obtained as weak subgraphs of the Cayley graphs of suitable E-groups where E is the set of edges of the graph to be covered (individually labelled as it were). While similar constructions have been used in [10, 11] and a precursor for special graphs in [5], we illustrate the key to the new generalisation in Proposition 5.1 below.

(2) Natural reduced products with N-acyclic $\mathbb{I}$-groupoids yield finite branched N-acyclic coverings of hypergraphs where $\mathbb{I}$ encodes the intersection pattern between hyperedges in the given hypergraph (cf. $\mathbb{I}$ in (3); see [12]).

(3) A new and more direct approach to finite branched N-acyclic coverings of hypergraphs (V, S) can be based on $\mathbb{I}$-products between a constraint graph $\mathbb{I} = (S, \mathsf{E})$ induced by the intersection graph of (V, S) and suitable E-groups that are not just N-acyclic but N-acyclic over $\mathbb{I}$; cf. Proposition 5.2 below.

Of these fundamental applications, (2) has been explored in stages in [11, 12, 13]. Application (1) is new in its strong form that involves the new notion of N-acyclicity of groups over a constraint graph $\mathbb{I}$. Application (3) similarly supersedes (2); also based on the new notion of N-acyclicity of groups over a constraint graph $\mathbb{I}$, it allows us to circumvents the use of groupoids in hypergraph coverings. Recall from Section 3.3 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 4 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).

For a finite simple graph $\mathbb{V} = (V, E)$ consider, as a set E of involutive generators for E-groups, the set of all edges $e = (v, v') \in E$, and as a constraint graph $\mathbb{I}$ the E-graph $\mathbb{I} = (V, (\{e\})_{e \in \mathsf{E}})$ ($\mathbb{V}$ with individually labelled edges). For any E-group $\mathbb{G}$ the natural projection

$$\begin{aligned} \pi \colon \mathbb{I} \otimes \mathbb{C}\mathbb{G} &\longrightarrow \mathbb{V} \\ (v, g) &\longmapsto v \end{aligned}$$

provides an unbranched covering of $\mathbb{V}$ by $\hat{\mathbb{V}} := \mathbb{I} \otimes \mathbb{C}\mathbb{G}$. Recall that $\mathbb{I} \otimes \mathbb{C}\mathbb{G}$ consists of components $\mathbb{C}\mathbb{G}_v$ for $v \in V$ (the vertex set of $\mathbb{I}$), which are pairwise isomorphic if $\mathbb{V}$ is connected, cf. Remark 3.11. In this case, the restriction of π to one of these connected components of $\hat{\mathbb{V}}$, each of which is a weak subgraph of $\mathbb{C}\mathbb{G}$, also is an unbranched covering.

Proposition 5.1. *Let $\mathbb{V} = (V, E)$ be a connected finite simple graph, E associated with its edge set E as above and $\mathbb{G}$ an E-group that is compatible with the E-graph $\mathbb{I} := (V, (\{e\})_{e\in\mathsf{E}})$. Then each connected component of the $\mathbb{I}$-product $\mathbb{I}\otimes\mathbb{C}\mathbb{G}$, $\hat{V} = \mathbb{C}\mathbb{G}_v$, which is realised as a weak subgraph of the Cayley graph $\mathbb{C}\mathbb{G}$ of $\mathbb{G}$, is an unbranched finite covering w.r.t. the natural projection $\pi\colon (v, g) \mapsto v$. This covering graph $\hat{\mathbb{V}}$ inherits the acyclicity properties of $\mathbb{C}\mathbb{G}$ in the sense of Remark 3.10: if $\mathbb{G}$ is N-acyclic over $\mathbb{I}$, then $\hat{\mathbb{V}}$ admits no cyclic configurations of up to N non-trivially overlapping α_i-connected components for subsets $\alpha_i \subseteq E$ (cf. Remark 3.10 for the precise statement).*

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

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

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

$$\bigcup_{s\in S} s \times \{g \in G \colon (s, g) \in \mathbb{I}\otimes\mathbb{C}\mathbb{G}\}$$

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

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

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

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

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

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

with covering projection

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

Proposition 5.2. *Let (V, S) be a finite hypergraph, $\mathbb{I} = (S, \mathsf{E})$ its intersection graph. If $\mathbb{G}$ is an E-group that is compatible with $\mathbb{I}$ then the hypergraph $\hat{\mathbb{V}} := \mathbb{V}\otimes\mathbb{C}\mathbb{G}$, which is based on the $\mathbb{I}$-product $\mathbb{I}\otimes\mathbb{C}\mathbb{G}$ of $\mathbb{I}$ with the Cayley graph $\mathbb{C}\mathbb{G}$ of $\mathbb{G}$, gives rise to a finite branched hypergraph covering $\pi\colon \hat{\mathbb{V}} \longrightarrow \mathbb{V}$. This covering hypergraph $\hat{\mathbb{V}}$ inherits the acyclicity properties of $\mathbb{C}\mathbb{G}$ in the following sense: if $\mathbb{G}$ is N-acyclic over $\mathbb{I}$, then every induced sub-hypergraph on up to N vertices is acyclic in the sense of classical hypergraph theory.*

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

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

It remains to argue for N-acyclicity of $\hat{\mathbb{V}}$ if $\mathbb{G}$ is chosen to be N-acyclic over $\mathbb{I}$. (Cf. Remark 3.3 for acyclicity in classical hypergraph terminology.) We show that in this situation the Gaifman graph of $\hat{\mathbb{V}}$ has no chordless cycles of lengths n with $3 < n \leqslant N$, nor cliques of size up to N that are not contained in any one of its hyperedges.

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

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

is an $\mathbb{I}$-coset cycle in $\mathbb{G}$, in the sense of Definition 3.8. Then $n > N$ follows from N-acyclicity of $\mathbb{G}$ over $\mathbb{I}$. Of the conditions in Definition 3.8, the first two are obvious for the given data: $s_{i+1} = h_i^{\mathbb{I}}(s_i)$ and $g_{i+1} \in g_i \mathbb{G}[\mathbb{I}, s_i, \alpha_i]$ for all $i \in \mathbb{Z}_n$. It remains to check that

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

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

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

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

as a candidate $\mathbb{I}$-coset cycle. We show that if this is not an $\mathbb{I}$-coset cycle, then the whole of m is contained in some hyperedge $[s, g] \in \hat{S}$. Again, the first two conditions on $\mathbb{I}$-coset cycles from Definition 3.8 are obvious for the given data:

$s_{i+1} = h_i^{\mathbb{I}}(s_i)$ and $g_{i+1} \in g_i\mathbb{G}[\mathbb{I}, s_i, \beta_i]$ since $\hat{v}_{i+1} \in [s_i, g_i] \cap [s_{i+1}, g_{i+1}]$. The intersection condition now is

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

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

References

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

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

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

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

[5] 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.

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

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

[8] I. Hodkinson and M. Otto. Finite conformal hypergraph covers and Gaifman cliques in finite structures. *Bulletin of Symbolic Logic*, 9, pp. 387–405, 2003.

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

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

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

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

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