

On Groupoids and Hypergraphs

Martin Otto
Department of Mathematics
Technische Universität Darmstadt

November 23, 2012

Abstract

We present a novel construction of finite groupoids whose Cayley graphs have large girth even w.r.t. a discounted distance measure that contracts arbitrarily long sequences of edges from the same colour class (subgroupoid), and only counts transitions between colour classes (cosets). These groupoids are employed towards a generic construction method for finite hypergraphs that realise specified overlap patterns and avoid small cyclic configurations. The constructions are based on reduced products with groupoids generated by the elementary local extension steps, and can be made to preserve the symmetries of the given overlap pattern. In particular, we obtain highly symmetric, finite hypergraph coverings without short cycles. The groupoids and their application in reduced products are sufficiently generic to be applicable to other constructions that are specified in terms of local glueing operations and require global finite closure.

Contents

1	Introduction	3
2	Groupoids	6
2.1	I-graphs	7
2.2	I-graphs and groupoids	9
2.3	Amalgamation of I-graphs	13
2.4	Avoiding some cycles	15
3	Hypergraphs, covers, and unfoldings	17
3.1	Hypergraphs and hypergraph covers	17
3.2	Covers by reduced products with groupoids	19
3.2.1	Chordality in reduced products	20
3.2.2	Conformality in reduced products	20
3.3	N-acyclic covers	21
3.4	Variations on unfoldings	21
3.4.1	Unfoldings w.r.t. link structure	21
3.4.2	Unfoldings of I-graphs	23
3.4.3	Realisations of overlap patterns	24
4	Symmetries	26
4.1	Groupoidal symmetries	26
4.2	Lifting structural symmetries	26
4.3	An application to Herwig's theorem	29
4.4	Regular finite N-acyclic hypergraphs	35

1 Introduction

Hypergraphs are structures consisting of vertices that form vertex clusters called hyperedges; formally the hyperedges are just subsets of the vertex set. The more familiar case of simple graphs is captured by the constraint that all hyperedges must have size 2 (2-uniform hypergraphs). Hypergraphs are the adequate combinatorial abstraction for various settings in which global structure is determined by local patches and their overlap pattern. Among the most important criteria of structural simplicity in hypergraphs are acyclicity constraints – a natural generalisation of graph acyclicity, which is qualitatively more complex.

Hypergraph covers aim to reproduce the overlap patterns between hyperedges of a given hypergraph in a covering hypergraph while smoothing out the overall behaviour, e.g., by achieving a higher degree of acyclicity. In graphs one would without ambiguity appeal to local pattern versus global structure: local structure manifests itself in the incidence degrees of edges in individual vertices, while global structure manifests itself, e.g., in the length of the shortest cycles, also called the girth of the graph. In hypergraphs the situation is more complicated because the boundary between local and global aspects is blurred since the transition from one hyperedge to the next typically preserves several vertices while exchanging others.

Looking at the example of the full 3-uniform hypergraph on a set of 4 vertices – also familiar as the boundary of the 3-simplex, or as the combinatorial structure formed by the faces of the tetrahedron – we see that every vertex is at the centre of and incident with a 3-cycle of those three faces that share this vertex. The transitions between these faces each involve the exchange of just one vertex while fixing a pair of vertices (those that form the edge in which these two faces are joined). Any natural process of finite unfolding will locally produce a $3n$ -cycle instead of the 3-cycle, still centred on a shared vertex, and thus finite covers cannot avoid cycles even locally. Moreover, even this local view informs us that non-trivial covers will affect incidence degrees and we have to deal with branched rather than unbranched covers: comparing the incidence degree of the pivotal vertex in a (locally around this vertex) k -fold unfolding to the incidence degrees in the original hypergraph, we see that each one of the original hyperedges is covered by k copies incident on a single shared vertex.

Very symmetric infinite coverings that are locally k -fold covers, hence locally finite, and have shortest (local) cycles of length $3k$, arise from regular isometric tessellations by equilateral triangles in the hyperbolic plane. One just has to use equilateral triangles of the right angular defect (i.e., hyperbolic area) to have exactly $3k$ such triangles incident in every point, each covering an angle of $2\pi/3k$. It seems far from obvious, however, how this elegant locally finite but infinite construction could be used towards the construction of finite covers with similar acyclicity and, if possible, symmetry properties.

Both of the perceived obstacles from the simple example, viz. the lack of a clear local-vs-global distinction and the necessity to consider branched coverings, generally seem to stand in the way of a straightforward construction of finite locally acyclic hypergraph covers. Neither of these obstacles arises in the special

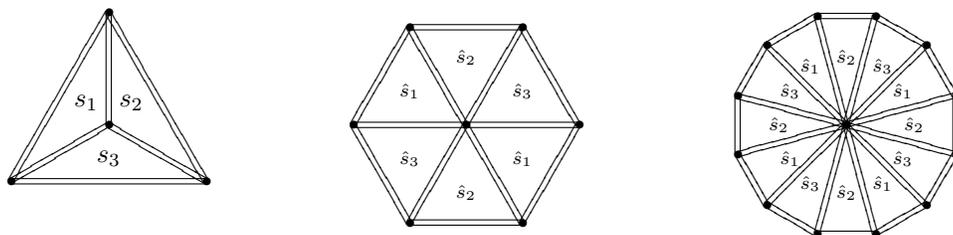


Figure 1: Local unfoldings of the tetrahedron, 2-fold and 4-fold.

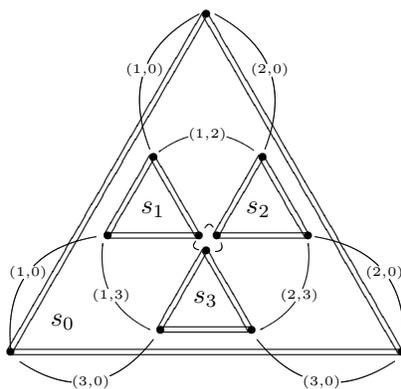


Figure 2: Overlap specification for the tetrahedron.

case of graphs, as the following result from [10] shows.

Proposition 1.1. *Every finite graph admits, for each $N \in \mathbb{N}$, a faithful (i.e., unbranched, degree-preserving) cover by a finite graph of girth greater than N (i.e., without cycles of length up to N).*

The uniform and canonical construction of such N -acyclic graph covers given in [10] is based on a natural product between the given graph with Cayley groups of large girth. The latter in turn can be obtained as subgroups of the symmetric groups of the vertex sets of suitably coloured finite acyclic graphs in an elegant combinatorial construction attributed to Biggs [4] in Alon’s survey [1].

Some of these ideas were successfully lifted and applied to the construction of hypergraph covers in [11]. The generalisation involved on the side of the combinatorial groups led to a uniform construction of Cayley groups that not only have large girth in the usual sense. Instead, they have large girth even w.r.t. to a reduced distance measure that measures the length of cycles in terms of the number of non-trivial transitions between cosets w.r.t. subgroups generated by different collections of generators. For an intuitive idea how this concern arises we may again look at the above example of the faces of the tetrahedron. There are two distinct sources of avoidable short cycles in its branched coverings: (a) ‘local cycles’ around a single pivot vertex, in the 1-neighbourhoods of a single vertex, and of length $3k$ in a locally k -fold unfolding; (b) ‘non-local cycles’ that enter and leave the 1-neighbourhoods of several distinct vertices. To account for the length of a cycle of type (b), the number of individual single-step transitions between faces around one of the visited pivotal vertices is typically irrelevant; what essentially matters is how often we move from one pivot to the next, and this corresponds to a transition between two subgroups (think of a transition between the stabiliser of one pivot and the next).

But nothing as simple as a (reduced) product between a hypergraph and even one of these ‘highly acyclic’ Cayley graphs will produce a covering by a finite N -acyclic hypergraph (to be defined properly below). The construction presented in [11] uses such Cayley groups only as one ingredient to achieve suitable hypergraph covers through an intricate local-to-global construction by induction on the width (the maximal size of hyperedges). More importantly, these further steps in the construction from [11], are no longer canonical or natural, in the sense, e.g., that they do not preserve symmetries of the given hypergraph. It also remains unclear in that construction which kinds of singularities and branching are unavoidable as opposed to artefacts due to non-canonical choices.

We now expand the amalgamation techniques that were explored for the combinatorial construction of highly acyclic Cayley graphs [11] from groups to groupoids, and obtain ‘Cayley groupoids’ that are highly acyclic in a similar sense.

It turns out that groupoids are a much better fit for the task of constructing hypergraph covers as well as for the construction of finite hypergraphs according to other specifications. The new notion of Cayley groupoids allows for the construction of covers by means of natural reduced products with these groupoids.

It is more canonical and supports covers of far greater genericity and symmetry than previously available. It also allows for substantial adaptations and generalisations. We address the covering problem and some variations and prove the following main theorems.

Theorem 1. Every finite hypergraph admits, for every $N \in \mathbb{N}$, a covering by a finite hypergraph that is N -acyclic, i.e., in which every induced sub-hypergraph of up to N vertices is acyclic. In addition, the covering hypergraph can be chosen to preserve all symmetries of the given hypergraph. See Section 3.1 for definitions and details, and especially Proposition 3.9.

Theorem 2. Every finite hypergraph admits finite unfoldings w.r.t. any specified link structure that stipulates (arbitrary and possibly multiple) overlaps between hyperedges that can be proper subsets of the overlaps realised in the given hypergraph. Such finite unfoldings can again be required to be N -acyclic and to preserve all symmetries of the given hypergraph and the specified link structure. See Section 3.4.1 and especially Proposition 3.13.

Theorem 3. Every abstract finite specification of an overlap pattern between disjoint sets, by means of partial matchings between them, admits a finite realisation by a finite hypergraph, which may again be required to be N -acyclic and to preserve all symmetries of the given specification. See Section 3.4.3 and Proposition 3.17.

In this first presentation of the new technique I emphasise the technical notions and key lemmas and relegate an analysis of further applications to further study. The body of this note is organised in three sections. The first of these, Section 2, introduces the notion of Cayley groupoids and presents a combinatorial method for their construction that is inspired by the work on Cayley groups in [11] and ultimately by Biggs [4]. The second main part, Section 3, presents hypergraphs and covers obtained as natural reduced products with Cayley groupoids and shows how the acyclicity criteria achieved for groupoids in Section 2 translate into degrees of acyclicity for finite hypergraph covers. Finally, in Section 4, the symmetries hypergraphs constructed as reduced products with Cayley groupoids are explored; as one further application, a variant of Herwig’s theorem on the extension of partial isomorphism is discussed.

2 Groupoids

In this section we develop a method to obtain groupoids from operations on coloured graphs. The basic idea is similar to the construction of Cayley groups as subgroups of the symmetric group of the vertex set of a graph, where the subgroup is generated by permutations induced by the graph structure and in particular by the edge colouring of the graphs in question. This method is useful for the construction of Cayley groups and associated homogeneous

graphs of large girth [1, 4]. In that case, one considers simple undirected graphs $H = (V, (R_e)_{e \in E})$ with edge colours $e \in E$ such that every vertex is incident with at most one edge of each colour. Then $e \in E$ induces a permutation π_e of the vertex set V , where π_e swaps the two vertices in every e -coloured edge. The $(\pi_e)_{e \in E}$ generate a subgroup of the group of all permutations of V . For suitable H , the Cayley graph induced by this group with generators $(\pi_e)_{e \in E}$ can be shown to have large girth (no short cycles, i.e., no short generator sequences that represent the identity). We here expand the underlying technique from groups to groupoids and lift it to a higher level of ‘large girth’. The second aspect is similar to the strengthening obtained in [11] for groups. The shift in focus from groups to groupoids is new here. Just as Cayley groups and their Cayley graphs, which are particularly homogeneous edge-coloured graphs, are extracted from group actions on given edge-coloured graphs in [11], we shall here construct groupoids and associated groupoidal Cayley graphs, which are edge- and vertex-colored graphs (I -graphs in the terminology introduced below), from given I -graphs. The generalisation from Cayley groups to the new *Cayley groupoids* requires conceptual changes and presents quite some additional technical challenges, but leads to objects that are better suited to hypergraph constructions than Cayley groups.

2.1 I-graphs

Let S be a finite set, and for $s \neq s' \in S$, $E[s, s']$ a finite set of edge colours for directed edges from s to s' . The family of edge colours $E = (E[s, s'])_{s \neq s'}$ is equipped with a converse operation, which maps the edge colour $e \in E[s, s']$ to the edge colour $e^{-1} \in E[s', s]$, as an involutive bijection between $E[s, s']$ and $E[s', s]$. We write $I = (S, E) = (S, (E[s, s']))$ for this graph template, which may be visualised as a *loop-free multi-graph*, in which several edges may link the same two vertices s and s' , and that is directed but closed under edge reversal, with distinct labels for the two directions of every edge.¹ It is often convenient to write E also for the disjoint union of the $E[s, s']$.

Definition 2.1. An *I-graph* is a finite directed edge- and vertex-coloured graph $H = (V, (V_s)_{s \in S}, (R_e)_{e \in E})$, whose vertex set V is partitioned into subsets V_s and in which, for $e \in E[s, s']$, the directed edge relation $R_e \subseteq V_s \times V_{s'}$ induces a partial bijection from V_s to $V_{s'}$ (i.e., R_e is a partial matching between V_s and $V_{s'}$), such that the $(R_e)_{e \in E}$ are compatible with converses: $R_{e^{-1}} = (R_e)^{-1}$. H is a *complete I-graph* if the R_e induce full rather than partial matchings.

Edges in R_e are also referred to as edges of colour e or just as e -edges. We may regard I -graphs as a restricted class of S -partite, E -coloured multi-graphs.

A process of *completion* is required to prepare arbitrary given I -graphs for the desired groupoidal operation.

¹A cleaner formalisation would be a two-sorted framework with a node set S and a set E of directed edges together with incidence maps $\iota_0, \iota_1 : E \rightarrow S$ specifying the source and target nodes of every edge; for our purposes, however, this imposes unnecessary burdens on notation.

If $H = (V, (V_s), (R_e))$ is an I -graph then the following produces a complete I -graph on the vertex set $V \times S$, with the partition induced by the natural projection, and with a natural embedding of H into $H \times I$:

$$\begin{aligned} H \times I &:= (V \times S, (V_s), (R_e)) \text{ where} \\ V_s &= V \times \{s\}, \text{ for } s \in S, \end{aligned}$$

and, for $e \in E[s, s']$,

$$\begin{aligned} R_e &= \{((v, s), (v', s')) : (v, v') \text{ an } e\text{-edge in } H\} \cup \\ &\quad \{((v', s), (v, s')) : (v, v') \text{ an } e\text{-edge in } H\} \cup \\ &\quad \{((v, s), (v, s')) : v \text{ not incident with an } e\text{-edge in } H\}, \end{aligned}$$

with embedding

$$\begin{aligned} \sigma: V &\longrightarrow V \times S \\ v &\longmapsto (v, s) \text{ for } v \in V_s. \end{aligned}$$

We note that this stipulation does indeed produce a complete I -graph: it is clear from the definition of the R_e , for $e \in E[s, s']$, that $R_e \subseteq V_s \times V_{s'}$ and that every vertex in V_s has an outgoing e -edge and every vertex in $V_{s'}$ an incoming e -edge; R_e also is a matching as required: either $v \in V$ is incident with an e -edge in H , which means that, for a unique $v' \in V$, one of (v, v') or (v', v) is an e -edge in H , and in both cases $((v, s), (v', s'))$ and $((v', s), (v, s'))$ become e -edges in $H \times I$; or v is not incident with an e -edge in H , and $((v, s), (v, s'))$ thus becomes the only outgoing e -edge from (v, s) as well as the only incoming e -edge at (v, s') . Also $R_{e^{-1}} = (R_e)^{-1}$ as required.

Observation 2.2. *If $H = (V, (V_s), (R_e))$ is a not necessarily complete I -graph, then $H \times I$ is a complete I -graph; the embedding*

$$\begin{aligned} \sigma: V &\longrightarrow V \times S \\ v &\longmapsto (v, s) \text{ for } v \in V_s \end{aligned}$$

embeds H isomorphically as an induced substructure into $H \times I$.

Proof. Note that the natural projection onto the first factor provides the inverse to σ on its image. Then $(v, v') \in R_e^H$ for $e \in E[s, s']$ implies that $v \in V_s$ and $v' \in V_{s'}$ and therefore that $(\sigma(v), \sigma(v')) = ((v, s), (v', s'))$ is an e -edge of $H \times I$. Conversely, let $(\sigma(v), \sigma(v')) = ((v, s), (v', s'))$ be an e -edge of $H \times I$; then $s \neq s'$ implies that $v \neq v'$ (as the V_s partition V), whence (v, v') must be an e -edge of H . \square

In the following we use, as a completion of H , the relevant connected component(s) of $H \times I$; i.e., the components into which H naturally embeds.

Definition 2.3. The *completion* \bar{H} of a not necessarily complete I -graph $H = (V, (V_s), (R_e))$ is the union of the connected components in $H \times I$ incident with the vertex set $\sigma(V) = \{(v, s) : v \in V_s\}$.

Identifying V with $\sigma(V) \subseteq H \times I$, we may regard \bar{H} as an extension of H .

Corollary 2.4. *For every I -graph H , the completion \bar{H} is a complete I -graph. Completion is compatible with disjoint unions: if $H = H_1 \dot{\cup} H_2$ is a disjoint union of I -graphs H_i , then $\bar{H} = \bar{H}_1 \dot{\cup} \bar{H}_2$. If H itself is complete, then $\bar{H} \simeq H$.*

Proof. The first claim is obvious: by definition of completeness, any union of connected components of a complete I -graph is itself complete.

For compatibility with disjoint unions observe that the connected component of the σ -image of H_1 in $H \times I$ is contained in the cartesian product of H_1 with S , as edges of $H \times I$ project onto edges of H or onto loops.

For the last claim observe that, for complete H , the vertex set of the isomorphic embedding $\sigma: H \rightarrow H \times I$ is closed under the edge relations R_e of $H \times I$: due to completeness of H , every vertex in $\sigma(V_s)$ is matched to precisely one vertex in $\sigma(V_{s'})$ for every $e \in E[s, s']$; it follows that no vertex in $\sigma(V)$ can have additional edges to nodes outside $\sigma(V)$ in $H \times I$. \square

Where this is convenient, we may always assume w.l.o.g. that I itself is connected, and often also treat connected components of the I -graphs H under consideration separately. Connected components w.r.t. to subsets of the edge colours E , however, require special treatment. If $\alpha = \alpha^{-1} \subseteq E$ we write I_α for the reduct of I to its α -edges – of course I_α may be disconnected even if I is connected. We regard the α -reducts of I -graphs (literally: their reducts to just those binary relations R_e for $e \in \alpha$) as I_α -graphs; the α -reduct of the I -graph H is denoted $H \upharpoonright \alpha$. Connected components of I -graphs w.r.t. α -edges (edges of colours $e \in \alpha$) will arise as typical I_α -graphs in some constructions below.

Lemma 2.5. *Let $\alpha = \alpha^{-1} \subseteq E$, and consider an I -graph H and its α -reduct $K = H \upharpoonright \alpha$, as well as their closures as I -graphs, \bar{H} and \bar{K} , and the closure of K as an I_α -graph, \bar{K}^α . Then*

$$\bar{K}^\alpha \simeq \bar{K} \upharpoonright \alpha \simeq \bar{H} \upharpoonright \alpha.$$

Proof. It suffices to note that, for $e \in \alpha$, the completions mentioned only introduce new e -edges (in the α -connected components of $\sigma(K)$ or $\sigma(H)$) where e -edges are missing; and those are introduced in the same manner regardless of whether the completion is w.r.t. α or all of E . \square

2.2 I -graphs and groupoids

We now obtain a groupoid operation on every complete I -graph H generated by the local bijections $\rho_e: V_s \rightarrow V_{s'}$ for $e \in E[s, s']$ as induced by the R_e . This essential step supports a new groupoidal analogue of the passage from coloured graphs to Cayley groups, and the duality between Cayley groups and their Cayley graphs, at the level of I -graphs: we find a similar relationship between Cayley groupoids and their Cayley graphs.

Definition 2.6. An S -groupoid is a structure $\mathbb{G} = (G, (G_{st})_{s,t \in S}, \cdot, (1_s)_{s \in S})$ whose domain G is partitioned into the sets G_{st} , with designated $1_s \in G_{ss}$ for $s \in S$ and a partial binary operation \cdot on G which is precisely defined on the union of the sets $G_{st} \times G_{tu}$, where it takes values in G_{su} , such that the following conditions are satisfied:

- (i) (associativity) for all $g \in G_{st}, h \in G_{tu}, k \in G_{uv}$: $g \cdot (h \cdot k) = (g \cdot h) \cdot k$.
- (ii) (neutral elements) for all $g \in G_{st}$: $g \cdot 1_t = g = 1_s \cdot g$.
- (iii) (inverses) for every $g \in G_{st}$ there is some $g^{-1} \in G_{ts}$ such that $g \cdot g^{-1} = 1_s$ and $g^{-1} \cdot g = 1_t$.

Recall that $I = (S, E)$ is a fixed directed and loop-free multi-graph with converse edges, which serves as the template for I -graphs. For simplicity E stands for both the family $E = (E[s, s'])_{s \neq s'}$ and the disjoint union over its members. In the following, the edge set E is used as the set of edge colours in I -graphs, and also associated with generators of groupoids. Recall also that E is endowed with the converse operation $e \mapsto e^{-1}$, which bijectively relates $E[s, s']$ and $E[s', s]$ for every pair $s \neq s'$ in S .

For $I = (S, E)$, we let E^* stand for the set of all labellings of directed paths in I . A typical element of E^* thus is of the form $w = e_1 \dots e_n$ where $n \in \mathbb{N}$ is its length, and for suitable $s_i \in S$, the edges are such that $e_i \in E[s_i, s_{i+1}]$ for $1 \leq i \leq n$. We admit the empty labellings of paths of length 0 at $s \in S$, and distinguish them by their location s as λ_s .² The set E^* is partitioned into subsets E_{st}^* , which, for $s, t \in S$, consist of the labellings of paths from s to t in I , so that in particular $\lambda_s \in E_{ss}^*$. If $w = e_1 \dots e_n \in E_{st}^*$ we write $w^{-1} := e_n^{-1} \dots e_1^{-1}$ for the *converse* in E_{ts}^* , which is obtained by reverse reading w and replacement of each edge label e by its converse e^{-1} . The set E^* carries a partially defined associative concatenation operation

$$(w, w') \in E_{st}^* \times E_{tu}^* \quad \mapsto \quad ww' \in E_{su}^*,$$

which has the empty words $\lambda_s \in E_{ss}^*$ as neutral elements. One may think of this structure as a groupoidal analogue of the familiar word monoids. For further reference, we denote it as the *free I -structure*

$$\mathfrak{F}^* = (E^*, (E_{st})_{s,t \in S}, \cdot, (\lambda_s)_{s \in S}, {}^{-1}),$$

and note that the converse operation $w \mapsto w^{-1}$ does not provide inverses, but will produce groupoidal inverses under homomorphisms to be considered.

Definition 2.7. The S -groupoid \mathbb{G} is *generated* by the family $(g_e)_{e \in E}$ if

- (i) for every $e \in E[s, s']$, $g_e \in G_{ss'}$ and $g_{e^{-1}} = (g_e)^{-1}$;
- (ii) for every $s, t \in S$, every $g \in G_{st}$ is represented by a product $\prod_{i=1}^n g_{e_i}$, for some path $e_1 \dots e_n \in E_{st}^*$.

²Note that for convenience we use the notation E^* , which usually stands for the set of all E -words, with a different meaning: firstly, E^* here only comprises E -words that arise as labellings of directed paths in I ; but secondly, we distinguish empty words $\lambda_s \in E^*$, one for every $s \in S$.

An S -groupoid \mathbb{G} that is generated by some family $(g_e)_{e \in E}$ is called an I -groupoid.³

In other words, an I -groupoid is a groupoid that is a homomorphic image of the free I -structure \mathfrak{E} , under the map

$$\begin{aligned} \mathbb{G}: \mathfrak{J}^* &\longrightarrow \mathbb{G} \\ w = e_1 \dots e_n \in E_{st}^* &\longmapsto w^{\mathbb{G}} = \prod_{i=1}^n g_{e_i} \in G_{st}. \end{aligned}$$

Note that, if I is connected, then an I -groupoid is also connected in the sense that any two groupoid elements are linked by a path of generators. Otherwise, for disconnected I , an I -groupoid breaks up into connected components that form separate groupoids, viz., one I' -groupoid for each connected component I' .

For a subset $\alpha = \alpha^{-1} \subseteq E$ that is closed under converse we denote by \mathbb{G}_α the sub-groupoid generated by $(g_e)_{e \in \alpha}$ within \mathbb{G} :

$$\mathbb{G}_\alpha = \mathbb{G} \upharpoonright \{w^{\mathbb{G}} : w \in \bigcup_{st} \alpha_{st}^*\} \text{ with generators } (g_e)_{e \in \alpha},$$

which, according to the above, may break up into separate and disjoint $I_{\alpha'}$ -groupoids for the disjoint connected components $I_{\alpha'}$ of I_α .

Consider a not necessarily complete I -graphs $H = (V, (V_s), (R_e))$. The partial bijections prescribed by the relations R_e , together with their compositions along paths in E^* , induce a structure of the same type as the free \mathfrak{J}^* , in fact a natural homomorphic image of \mathfrak{J}^* .

For $e \in E[s, s']$, let $\rho_e = \rho_e^H$ be the partial bijection between V_s and $V_{s'}$ induced by $R_e \subseteq V_s \times V_{s'}$ (we drop the superscripts H where this causes no confusion). For $w \in E_{st}^*$, define ρ_w as the partial bijection from V_s to V_t induced by the composition of the maps ρ_{e_i} along the path $w = e_1 \dots e_n$; in relational terminology, the graph of ρ_w is the relational composition of the R_{e_i} . For $w \in E_{st}^*$,

$$\rho_w (= \rho_w^H): V_s \longrightarrow V_t$$

is a partial bijection, possibly empty. In this manner we obtain a homomorphic image of the free I -structure $\mathfrak{J}^* = (E^*, (E_{st}) \cdot, (\lambda_s), {}^{-1})$.

Concatenation maps to (partial) composition:

$$\rho_{ww'} = \rho_{w'} \circ \rho_w \text{ (composition of partial maps)}$$

wherever defined, i.e., for $w \in E_{st}^*$, $w' \in E_{t'u}^*$ so that $ww' \in E_{su}^*$.

The converse operation $w \mapsto w^{-1}$ maps to the inversion $\rho_w \mapsto \rho_{w^{-1}} = (\rho_w)^{-1}$ of partial maps

$$(\rho_w)^{-1} = \rho_{w^{-1}}.$$

In general this is not a proper groupoidal inverse, since, for $e \in E[s, s']$ say, $\text{dom}((\rho_e^H)^{-1} \circ (\rho_e^H)) = \text{dom}(\rho_e^H) \subseteq V_s$ may be a proper subset. In fact it is the

³It will often make sense to identify the generator g_e with e itself, and we shall often also speak of groupoids generated by the family $(e)_{e \in E}$.

crucial distinguishing feature of complete I -graphs that we obtain a groupoidal inverse.

In general, i.e., for an arbitrary I -graph H , the homomorphic image of the free I -structure $\mathfrak{E} = (E^*, (E_{st}), \cdot, (\lambda_s), {}^{-1})$ under the map

$$\begin{aligned} \rho^H: \mathfrak{J}^* &\longrightarrow \{ \rho: \rho \text{ a partial bijection of } V \} \\ w = e_1 \dots e_n &\longmapsto \rho_w = \prod_{i=1}^n \rho_{e_i} \end{aligned}$$

produces a structure without groupoidal inverses. For a complete I -graph H , however, $\rho_{w^{-1}} \circ \rho_w = (\rho_w)^{-1} \circ \rho_w = \text{id}_{V_s}$ for any $w \in E_{st}^*$, and the image structure obtained in this manner is a groupoid $\mathbb{G} =: \text{cym}(H)$:

$$\rho^H: \mathfrak{J}^* \longrightarrow \mathbb{G} =: \text{cym}(H),$$

where

$$\begin{aligned} \text{cym}(H) = \mathbb{G} &= (G, (G_{st})_{s,t \in S}, \cdot, (1_s)_{s \in S}), \\ G_{st} &= \{ \rho_w: w \in E_{st}^* \}. \end{aligned}$$

The groupoid operation \cdot is the one imposed by the natural composition structure between members of these sorts:

$$\begin{aligned} \cdot: \bigcup_{s,t,u} G_{st} \times G_{tu} &\longrightarrow G \\ (\rho_w, \rho_{w'}) \in G_{st} \times G_{tu} &\longmapsto \rho_w \cdot \rho_{w'} := \rho_{ww'} \in G_{su}. \end{aligned}$$

For $s \in S$, the identity $1_s := \text{id}_{V_s}$ is the neutral element of sort G_{ss} , induced as $1_s = \rho_\lambda$ by the empty word $\lambda \in E_{ss}^*$.

It is clear from the discussion above that there is a natural groupoidal inverse

$$\begin{aligned} {}^{-1}: G &\longrightarrow G \\ \rho_w \in G_{st} &\longmapsto (\rho_w)^{-1} := \rho_{w^{-1}} \in G_{ts} \end{aligned}$$

as $\rho_{w^{-1}}$ is the full inverse $(\rho_w)^{-1}: V_t \rightarrow V_s$ of the bijection $\rho_w: V_s \rightarrow V_t$.

Definition 2.8. For a complete I -graph H we let $\text{cym}(H)$ be the groupoid abstracted from H according to the above stipulations. We consider $\text{cym}(H)$ as an I -groupoid generated by $(\rho_e)_{e \in E}$.

It is easy to check that $\text{cym}(H)$ is an I -groupoid with generators $(\rho_e)_{e \in E}$ according to Definition 2.7. We turn to the analogue, for I -groupoids, of the notion of Cayley graph.

Definition 2.9. Let $\mathbb{G} = (G, (G_{st}), \cdot, (1_s))$ be an I -groupoid generated by $(g_e)_{e \in E}$. The *Cayley graph* of \mathbb{G} (w.r.t. these generators) is the complete I -graph $(V, (V_s), (R_e))$ where $V = G$,

$$V_s = G_{*s} := \bigcup_t G_{ts},$$

and

$$R_e = \{(g, g \cdot e): g \in V_s\} \text{ for } e \in E[s, s'].$$

One checks that this stipulation indeed specifies a complete I -graph, and in particular that really $R_e \subseteq V_s \times V_{s'}$ for $e \in E[s, s']$.

Lemma 2.10. *The I -groupoid induced by the Cayley graph of \mathbb{G} is isomorphic to \mathbb{G} .*

Proof. Consider a generator ρ_e of the I -groupoid induced by the Cayley graph of \mathbb{G} . For $e = (s, s')$ this is the bijection

$$\begin{aligned} \rho_e: V_s = G_{*s} &\longrightarrow V_{s'} = G_{*s'} \\ g &\longmapsto g \cdot g_e, \end{aligned}$$

so that ρ_e operates as right multiplication by generator g_e (exactly where defined). Since the $(\rho_e)_{e \in E}$ generate the groupoid induced by the Cayley graph of \mathbb{G} , it suffices to show that groupoid products of the ρ_e (compositions) and the groupoid products of the g_e in \mathbb{G} satisfy the same equations, which is obvious from the correspondence just established. E.g., if $\prod_i g_{e_i} = 1_s$ in \mathbb{G} , then, for the corresponding $w = e_1 \dots e_n$, we have that $\rho_w: V_s \rightarrow V_s$, where $V_s = G_{*s}$, maps g to $g \cdot \prod_i g_{e_i} = g \cdot 1_s = g$ for all $g \in V_s = G_{*s}$, whence $\rho_w = \text{id}_{V_s}$ as desired. \square

If we identify I -groupoids with their Cayley graphs (which are complete I -graphs), we thus find that the generic process of obtaining I -groupoids from complete I -graphs trivially reproduces the given I -groupoid when applied to such. We extend the passage from I -graphs to I -groupoids to the setting of not necessarily complete I -graphs by combining it with the completion \bar{H} of H in $H \times I$.

Definition 2.11. For a not necessarily complete I -graph H , we let the induced I -groupoid $\text{cym}(H)$ be the I -groupoid $\text{cym}(\bar{H})$ induced (according to Definition 2.8) by the completion of H (in the sense of Definition 2.3).

2.3 Amalgamation of I -graphs

Consider two sub-groupoids \mathbb{G}_α and \mathbb{G}_β of an I -groupoid \mathbb{G} with generators $e \in E$, where $\alpha = \alpha^{-1}, \beta = \beta^{-1} \subseteq E$ are closed under converse. We write $\mathbb{G}_{\alpha\beta}$ for $\mathbb{G}_{\alpha \cap \beta}$ and note that $\alpha \cap \beta$ is automatically closed under converse.

For $g \in G_{*s}$ (a vertex of colour s in the Cayley graph) we may think of the connected component of g in the reduct of the Cayley graph of \mathbb{G} to those R_e with $e \in \alpha$ as the \mathbb{G}_α -coset at g :

$$g\mathbb{G}_\alpha = \{g \cdot w^{\mathbb{G}} : w \in \bigcup_t \alpha_{st}^*\} \subseteq G.$$

If I_α is connected, then $g\mathbb{G}_\alpha$, as a weak subgraph of (the Cayley graph of) \mathbb{G} , carries the structure of a complete I_α -graph. If I_α consists of disjoint connected components, then $g\mathbb{G}_\alpha$ really produces the coset w.r.t. $\mathbb{G}_{\alpha'}$ where α' is the edge set of the connected component of s in I_α . In any case, this I_α -graph is isomorphic to the connected component of 1_s in the Cayley graph of \mathbb{G}_α .

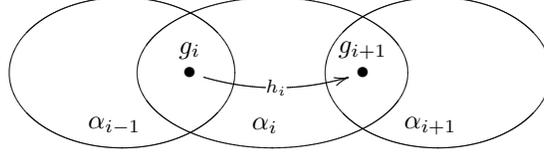


Figure 3: Amalgamation: overlap between cosets.

Suppose the I_α -graph H_α and the I_β -graph H_β are isomorphic to the Cayley graphs of sub-groupoids \mathbb{G}_α and \mathbb{G}_β , respectively. If $v_1 \in H_\alpha$ and $v_2 \in H_\beta$ are vertices of the same colour $s \in S$ (i.e., in terms of the isomorphisms with \mathbb{G}_α and \mathbb{G}_β , they correspond to group elements $g_1 \in \mathbb{G}_\alpha$ and $g_2 \in \mathbb{G}_\beta$ from the same G_{*s}), then the connected components w.r.t. edge colours in $\alpha \cap \beta$ of v_1 in H_α and of v_2 in H_β are related by a unique isomorphism (unique as an isomorphism between the weak subgraphs formed by the $(\alpha \cap \beta)$ -components). We define the amalgamation of (H_α, v_1) and (H_β, v_2) (with reference vertices v_1 and v_2 of the same colour s) to be the result of identifying the vertices in these two connected components in accordance with this unique isomorphism, and keeping everything else disjoint. It is convenient to speak of (the Cayley graphs of) the sub-groupoids \mathbb{G}_α as the constituents of such amalgams, but we keep in mind that we treat them as abstract I -graphs and not as embedded into \mathbb{G} – just locally, in the connected components of g_1 and g_2 , i.e. in $g_1\mathbb{G}_{\alpha\beta} \simeq g_2\mathbb{G}_{\alpha\beta}$, the structure of the amalgam is the same as the induced structure of $g\mathbb{G}_{\alpha\beta} \subseteq \mathbb{G}$ in \mathbb{G} for any $g \in V_s = G_{*s} \subseteq G$.

Let, in this sense,

$$(\mathbb{G}_{\alpha_1}, g_1) \oplus_s (\mathbb{G}_{\alpha_2}, g_2)$$

stand for the result of the amalgamation of the Cayley graphs of the two sub-groupoids \mathbb{G}_{α_i} in the vertices $g_i \in V_s \subseteq \mathbb{G}_{\alpha_i}$. Note that $(\mathbb{G}_{\alpha_1}, g_1) \oplus_s (\mathbb{G}_{\alpha_2}, g_2)$ is generally not a complete I -graph (or I_{α_i} -graph for either i) but satisfies the completeness requirement for edges $e \in \alpha_1 \cap \alpha_2$.

Let $(\mathbb{G}_{\alpha_i}, g_i, h_i, s_i)_{1 \leq i \leq N}$ be a sequence of sub-groupoids with distinguished elements and vertex colours as indicated, and such that, for $1 < i < N$,

$$(\dagger) \quad \begin{cases} g_i \in (G_{\alpha_i})_{*s_i} \subseteq \mathbb{G}_{\alpha_i} \\ h_i \in (G_{\alpha_i})_{s_i s_{i+1}} \subseteq \mathbb{G}_{\alpha_i} \\ g_i \mathbb{G}_{\alpha_{i-1} \alpha_i} \cap g_i h_i \mathbb{G}_{\alpha_i \alpha_{i+1}} = \emptyset \quad \text{as cosets in } \mathbb{G} \text{ (within } g_i \mathbb{G}_{\alpha_i} \text{)}. \end{cases}$$

Then the pairwise amalgams

$$(\mathbb{G}_{\alpha_i}, g_i h_i) \oplus_{s_i} (\mathbb{G}_{\alpha_{i+1}}, g_{i+1})$$

are individually well-defined and, due to the last requirement in (\dagger) , do not interfere. Together they produce a connected I -graph

$$H := \bigoplus_{i=1}^N (\mathbb{G}_{\alpha_i}, g_i, h_i, s_i).$$

Condition (†) is important to ensure that the resulting structure is an I -graph, since otherwise, if $\alpha_{i-1} \cap \alpha_{i+1} \neq \emptyset$, an element of the overlap would inherit new e -edges from both $\mathbb{G}_{\alpha_{i-1}}$ and from $\mathbb{G}_{\alpha_{i+1}}$, for $e \in (\alpha_{i-1} \cap \alpha_{i+1}) \setminus \alpha_i$. We call an amalgam produced in this fashion a *chain of sub-groupoids* \mathbb{G}_{α_i} of length N .

2.4 Avoiding some cycles

Definition 2.12. A *coloured cycle of length n* in an I -groupoid with generators $e \in E$ is a sequence $(g_i)_{i \in \mathbb{Z}_n}$ of groupoid elements g_i (cyclically indexed) together with a sequence of *colour sets* $\alpha_i = \alpha_i^{-1} \subseteq E$ such that

$$h_i := g_i^{-1} \cdot g_{i+1} \in \mathbb{G}_{\alpha_i}$$

and

$$g_i \mathbb{G}_{\alpha_i \alpha_{i-1}} \cap g_{i+1} \mathbb{G}_{\alpha_i \alpha_{i+1}} = \emptyset.$$

Definition 2.13. An I -groupoid with generators $e \in E$ is N -acyclic if it does not have coloured cycles of length up to N .

The following definition of compatibility captures the idea that some I -groupoid \mathbb{G} is at least as discriminating as the I -groupoid $\text{cym}(H)$ induced by the I -graph H . As indicated in the discussion following the definition, compatibility of \mathbb{G} with H says that $\text{cym}(H)$ is a homomorphic image of \mathbb{G} (in the sense of a homomorphism of I -groupoids, which respects the generators and is therefore compatible with the homomorphism from the free I -structure \mathfrak{J} onto $\text{cym}(H)$ and \mathbb{G}). It also means that $\mathbb{G} = \text{cym}(\mathbb{G}) = \text{cym}(\mathbb{G} \dot{\cup} H)$ – and in this role, compatibility of sub-groupoids \mathbb{G}_α with certain H will serve as a guarantee for the preservation of these (sub-)groupoids in construction steps that render the overall \mathbb{G} more discriminating.

Definition 2.14. For an I -groupoid \mathbb{G} and an I -graph H we say that \mathbb{G} is *compatible with H* if, for every $s \in S$ and $w \in E_{ss}^*$,

$$w^{\mathbb{G}} = 1_s \implies \rho_w = \text{id}_{V_s} = 1_s \text{ in } \text{cym}(H).$$

The condition of compatibility is such that the natural homomorphisms for the free \mathfrak{J}^* onto \mathbb{G} and onto $\text{cym}(H)$ induce a homomorphism from \mathbb{G} onto $\text{cym}(H)$, as in this commuting diagram:

$$\begin{array}{ccc} \mathfrak{J}^* & & \\ \downarrow \cdot_{\mathbb{G}} & \searrow \rho^H & \\ \mathbb{G} & \xrightarrow{\text{hom}} & \text{cym}(H) \end{array}$$

Note that, by definition, $\text{cym}(H)$ is compatible with H and \bar{H} and, by Lemma 2.10, with its own Cayley graph.

Remark 2.15. *If K and H are any I -graphs, then $\text{cym}(H \dot{\cup} K)$ is compatible with K , \bar{K} and with the Cayley graph of $\text{cym}(K)$.*

The following holds the key to avoiding short coloured cycles.

Lemma 2.16. *Let \mathbb{G} be an I -groupoid with generators $e \in E$, $k \in \mathbb{N}$, and assume that, for every $\alpha = \alpha^{-1} \subseteq E$ with $|\alpha| < 2k$, the sub-groupoid \mathbb{G}_α is compatible with chains of groupoids $\mathbb{G}_{\alpha\beta_i}$ up to length N , for any choice of subsets $\beta_i = \beta_i^{-1} \subseteq E$. Then there is a finite I -groupoid \mathbb{G}^* with the same generators s.t.*

- (i) *for every $\alpha = \alpha^{-1} \subseteq E$ with $|\alpha| < 2k$, $\mathbb{G}_\alpha^* \simeq \mathbb{G}_\alpha$, and*
- (ii) *for all $\alpha = \alpha^{-1} \subseteq E$ with $|\alpha| \leq 2k$, the sub-groupoid \mathbb{G}_α^* is compatible with chains $\mathbb{G}_{\alpha\beta_i}^*$ up to length N .*

Note that the closures of the amalgams that play a role in the lemma's compatibility assertions may be read either w.r.t. the respective $\alpha \subseteq E$ (thinking of \mathbb{G}_α as an I_α -groupoid) or w.r.t. all of E (with reference to \mathbb{G}_α as a sub-groupoid of the I -groupoid \mathbb{G}); by Lemma 2.5, the two readings are equivalent.

It will be important later that compatibility of \mathbb{G}_α^* with chains as in (ii) makes sure that \mathbb{G}_α^* cannot have coloured cycles (with colours $\alpha \cap \beta_i$ of length up to N , because every such cycle in the Cayley group $\mathbb{G}_\alpha^* = \text{cym}(\mathbb{G}_\alpha^*)$ would have to be a cycle also in the Cayley group induced by any such chain, including the one prescribed by the proposed coloured cycle itself (cf. Proposition 2.17 below).

Proof (of the lemma). We construct \mathbb{G}^* as $\mathbb{G}^* := \text{cym}(H)$ for an I -graph $H = \mathbb{G} \dot{\cup} K$ consisting of the disjoint union of (the Cayley graph of) \mathbb{G} and (the closures of) certain chains of sub-groupoids of G .

Specifically, we let K be the disjoint union of all amalgamation chains of length up to N of the form

$$\bigoplus_{i=1}^m (\mathbb{G}_{\alpha\beta_i}, g_i, h_i, s_i)$$

for $\alpha = \alpha^{-1}, \beta_i = \beta_i^{-1} \subseteq E$, $1 \leq i \leq m \leq N$, where $|\alpha| \leq 2k$.

By construction and Remark 2.15, $\mathbb{G}^* = \text{cym}(\mathbb{G} \cup K)$ is compatible with chains $\mathbb{G}_{\alpha\beta_i}$ of the required format; together with (i) this implies (ii), i.e., that \mathbb{G}^* is compatible with corresponding chains of $\mathbb{G}_{\alpha\beta_i}^*$: either the chain in question has only components $\mathbb{G}_{\alpha\beta}^*$ with $|\alpha \cap \beta| < 2k$ so that, by (i), $\mathbb{G}_{\alpha\beta}^* \simeq \mathbb{G}_{\alpha\beta}$; or there is some component $\mathbb{G}_{\alpha\beta}^*$ with $|\alpha \cap \beta| = 2k$, which implies that $\alpha = \beta \cap \alpha$ and by (\dagger), the merged chain is isomorphic to \mathbb{G}_α^* , thus trivialising the compatibility claim.

For (i), it suffices to show that, for $|\alpha'| < 2k$, $\mathbb{G}_{\alpha'}$ is compatible with each connected component of K . (That \mathbb{G}^* is compatible with $\mathbb{G}_{\alpha'}$ is clear since $\mathbb{G}_{\alpha'}$ is itself a component of K and hence of $\mathbb{G} \cup K$.) Compatibility with \mathbb{G} is obvious. Consider then a component of the form $\bigoplus_{i=1}^m (\mathbb{G}_{\alpha\beta_i}, g_i, h_i, s_i)$. Its α' -connected components are obtained as merged chains of components of the form $\mathbb{G}_{\alpha'\alpha\beta}$, cf. Lemma 2.5. Since $|\alpha'| < 2k$, the assumptions of the lemma imply compatibility of $\mathbb{G}_{\alpha'}$ with any such component. It follows that $\mathbb{G}^* = \text{cym}(\mathbb{G} \cup K)$

is compatible with all $\mathbb{G}_{\alpha'}$ for $|\alpha'| < 2k$, and thus $\mathbb{G}_{\alpha'}^* \simeq G_{\alpha'}$ for $|\alpha'| < 2k$. Compare Remark 2.15. \square

Proposition 2.17. *For every directed loop-free multi-graph $I = (S, E)$ with converse edges e^{-1} and $N \in \mathbb{N}$ there are finite N -acyclic I -groupoids with generators $e \in E$.*

Proof. Inductively apply Lemma 2.16 and note that the assumptions of the lemma are trivial for $k = 1$, because the trivial sub-groupoid generated by \emptyset , which just consists of the isolated neutral elements 1_s , is compatible with any I -graph. In each step as stated in the lemma, compatibility with corresponding chains implies that \mathbb{G}^* cannot have coloured cycles of length up to N with colour sequences of the form $\alpha\beta_i$ where $|\alpha| \leq 2k$. For $2k = |E|$, this rules out all coloured cycles of length up to N . \square

Observation 2.18. *For any 2-acyclic I -groupoid \mathbb{G} and any subsets $\alpha = \alpha^{-1}, \beta = \beta^{-1} \subseteq E$, with associated sub-groupoids $\mathbb{G}_\alpha, \mathbb{G}_\beta$ and $\mathbb{G}_{\alpha\beta}$:*

$$\mathbb{G}_\alpha \cap \mathbb{G}_\beta = \mathbb{G}_{\alpha\beta}.$$

Proof. Just the inclusion $\mathbb{G}_\alpha \cap \mathbb{G}_\beta \subseteq \mathbb{G}_{\alpha\beta}$ needs attention. Let $h \in \mathbb{G}_\alpha \cap \mathbb{G}_\beta$, i.e., $h = w^{\mathbb{G}} = (w')^{\mathbb{G}}$ for some $w \in \alpha_{st}^*$ and $w' \in \beta_{st}^*$. Let $g_0 \in G_{*s}$ and put $g_1 := g_0 \cdot h \in G_{*t}$. Then g_0, g_1 with $h_0 = g_0^{-1} \cdot g_1 = h \in \mathbb{G}_\alpha$ and $h_1 = g_1^{-1} \cdot g_0 = h^{-1} \in \mathbb{G}_\beta$, form a 2-cycle with colouring $\alpha_0 := \alpha, \alpha_1 := \beta$, unless the colouring condition

$$g_0 \mathbb{G}_{\alpha\beta} \cap g_1 \mathbb{G}_{\alpha\beta} = \emptyset$$

of Definition 2.12 is violated. So there must be some $k \in g_0 \mathbb{G}_{\alpha\beta} \cap g_1 \mathbb{G}_{\alpha\beta}$, which shows that $h = (g_0^{-1} \cdot k) \cdot (g_1^{-1} \cdot k)^{-1} \in \mathbb{G}_{\alpha\beta}$ as claimed. \square

3 Hypergraphs, covers, and unfoldings

3.1 Hypergraphs and hypergraph covers

A hypergraph is a structure $\mathfrak{A} = (A, S)$ where $S \subseteq \mathcal{P}(A)$ is called the set of hyperedges of \mathfrak{A} , A its vertex set.

Definition 3.1. With a hypergraph $\mathfrak{A} = (A, S)$ we associate

- (i) its *Gaifman graph* $G(\mathfrak{A}) = (A, G(S))$ where $G(S)$ is the simple undirected edge relation that links $a \neq a'$ in A if $a, a' \in s$ for some $s \in S$.
- (ii) its *intersection graph* $I(\mathfrak{A}) = (S, E)$ where $E = \{(s, s') : s \neq s', s \cap s' \neq \emptyset\}$.

The following criterion of hypergraph acyclicity is the natural and strongest notion of acyclicity (sometimes called α -acyclicity), cf., e.g., [3, 2]. It is in close correspondence with the algorithmically crucial notion of tree-decomposability (viz., existence of a tree-decomposition with hyperedges as bags) and with natural combinatorial notions of triangulation.

Definition 3.2. A finite hypergraph $\mathfrak{A} = (A, S)$ is *acyclic* if it is conformal and chordal:

- (i) conformality: every clique in the Gaifman graph $G(\mathfrak{A})$ is contained in some hyperedge $s \in S$;
- (ii) chordality: every cycle in the Gaifman graph $G(\mathfrak{A})$ of length greater than 3 has a chord.

For $N \geq 3$, $\mathfrak{A} = (A, S)$ is *N -acyclic* if it is N -conformal and N -chordal:

- (iii) N -conformality: every clique in the Gaifman graph $G(\mathfrak{A})$ of size up to N is contained in some hyperedge $s \in S$;
- (iv) N -chordality: every cycle in the Gaifman graph $G(\mathfrak{A})$ of length greater than 3 and up to N has a chord.

It would also be natural to define N -acyclicity in terms of N -chordality and full conformality. In our applications of these notions, however, N -conformality for sufficiently large N , viz. for N greater than the width of the hypergraph in question, implies outright conformality.

The induced sub-hypergraph $\mathfrak{A} \upharpoonright A_0$ of a hypergraph $\mathfrak{A} = (A, S)$ is the hypergraph on A_0 with hyperedge set $S \upharpoonright A_0 := \{s \cap A_0 : s \in S\}$. If \mathfrak{A} is N -acyclic, then every induced sub-hypergraph $\mathfrak{A} \upharpoonright A_0$ on subsets $A_0 \subseteq A$ of size up to N is acyclic.

A hypergraph homomorphism is a map $h: \mathfrak{A} \rightarrow \mathfrak{B}$ between hypergraphs $\mathfrak{A} = (A, S)$ and \mathfrak{B} such that, for every $s \in S$, $h \upharpoonright s$ is a bijection between the hyperedge s and some hyperedge $h(s)$ of \mathfrak{B} .

Definition 3.3. A hypergraph homomorphism $h: \hat{\mathfrak{A}} \rightarrow \mathfrak{A}$ between the hypergraphs $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ and $\mathfrak{A} = (A, S)$ is a *hypergraph cover* (of \mathfrak{A} by $\hat{\mathfrak{A}}$) if it satisfies the back-property w.r.t. hyperedges: for every $h(\hat{s}) = s \in S$ and $s' \in S$ there is some $\hat{s}' \in \hat{S}$ such that $h(\hat{s}') = s'$ and $h(\hat{s} \cap \hat{s}') = s \cap s'$.

Lemma 3.4 ([8]). *Every finite hypergraph admits a cover by a finite conformal hypergraph.*

The following observation is not going to be used in our present approach, but serves as yet another indicator that chordality, not conformality, poses the real challenge.

Observation 3.5. *Every hypergraph cover $h: \hat{\mathfrak{A}} \rightarrow \mathfrak{A}$ of a conformal hypergraph $\mathfrak{A} = (A, S)$ by some hypergraph $\hat{A} = (\hat{A}, \hat{S})$ gives rise to a cover by a conformal hypergraph $\hat{\mathfrak{A}}'$ obtained from $\hat{\mathfrak{A}}$ by extending \hat{S} to*

$$\hat{S}' := \hat{S} \cup \{\hat{t} : \hat{t} \subseteq \hat{A} \text{ a clique in } G(\hat{\mathfrak{A}})\}.$$

Moreover, $G(\hat{\mathfrak{A}}') = G(\hat{\mathfrak{A}})$ and, if $\hat{\mathfrak{A}}$ is (N -)chordal, then so is $\hat{\mathfrak{A}}'$.

Proof. First note that h is injective in restriction to any clique of $G(\hat{\mathfrak{A}})$ since it is injective in restriction to each hyperedge of $\hat{\mathfrak{A}}$. Conformality of $\hat{\mathfrak{A}}'$ and $G(\hat{\mathfrak{A}}') = G(\hat{\mathfrak{A}})$ are obvious from the definition. Thus conformality is gained, and (N -)chordality is preserved in the passage from $\hat{\mathfrak{A}}$ to $\hat{\mathfrak{A}}'$ since $G(\hat{\mathfrak{A}}') = G(\hat{\mathfrak{A}})$. \square

Remark 3.6. *The construction for Lemma 3.4 is canonical as is the extension step of Observation 3.5: they are compatible with automorphisms of the given hypergraph or hypergraph cover, and they are symmetric and preserve symmetries within fibres of the cover.*

3.2 Covers by reduced products with groupoids

Let $\mathfrak{A} = (A, S)$ be a finite hypergraph with intersection graph $I := I(\mathfrak{A}) = (S, E)$. This I is a directed loop-free graph with converse edges, rather than a multi-graph, since every $E[s, s']$ is either empty (if $s \cap s' = \emptyset$) or the singleton set $E[s, s'] = \{(s, s')\}$ (if $s \cap s' \neq \emptyset$). We treat this as a special case of a multi-graph.

Let \mathbb{G} be an I -groupoid with generators $e \in E$. For $a \in A$ we let \mathbb{G}_a denote the sub-groupoid of \mathbb{G} generated by $E_a \subseteq E$ where

$$E_a := \{(s, s') \in E : a \in s \cap s'\}.$$

We construct a natural hypergraph cover of \mathfrak{A} by a reduced product of \mathfrak{A} with \mathbb{G} :

$$\pi: \mathfrak{A} \otimes \mathbb{G} \longrightarrow \mathfrak{A},$$

where $\mathfrak{A} \otimes \mathbb{G} = (\hat{A}, \hat{S})$ is obtained as follows.

The vertex set \hat{A} is the quotient of the disjoint union of hyperedges s of \mathfrak{A} tagged by groupoid elements $g \in G_{*s}$,

$$\bigcup_{s \in S, g \in G_{*s}} \{g\} \times \{s\} \times s \quad (4)$$

w.r.t. the equivalence relation induced by identifications

$$(g, s, a) \approx (ge, s', a) \quad \text{for } e = (s, s') \in E_a.$$

We note that (g_1, s_1, a) is identified with (g_2, s_2, a) in this quotient if, and only if, there is a path w from s_1 to s_2 in I consisting of edges $e = (s, s') \in E_a$, for which $a \in s \cap s'$, and such that $g_2 = g_1 \cdot w^{\mathbb{G}}$.

We think of the generators $e = (s, s') \in E_a$ as preserving the vertex a in passage from $a \in s$ to $a \in s'$. I.e., we think of the g -tagged copy of s and the g' -tagged copy of s' , for $g' = ge$, as glued in their overlap $s \cap s'$.

Let us denote the equivalence class of a triple (g, s, a) as $[g, s, a]$. Then the hyperedges of $\mathfrak{A} \otimes \mathbb{G} = (\hat{A}, \hat{S})$ are the subsets represented by the natural copies of hyperedges $s \in S$:

$$\hat{S} = \{[g, s] : s \in S, g \in G_{*s}\} \text{ where } [g, s] := \{[g, s, a] : a \in s\} \subseteq \hat{A}.$$

The cover homomorphism π is the natural projection

$$\pi: [g, s, a] \longmapsto a.$$

It is easy to see that $\pi: \hat{\mathfrak{A}} \rightarrow \mathfrak{A}$ is indeed a hypergraph cover.

⁴Just for convenience, we add a tag to indicate membership of vertices $a \in s$ in s , although this is really redundant: s is already determined by g as the unique s for which $g \in G_{*s}$.

3.2.1 Chordality in reduced products

Lemma 3.7. *Let \mathfrak{A} be a hypergraph with intersection graph $I(\mathfrak{A}) =: I$, \mathbb{G} an N -acyclic I -groupoid. Then $\mathfrak{A} \otimes \mathbb{G}$ is N -chordal.*

Proof. Suppose that $([g_i, s_i, a_i])_{i \in \mathbb{Z}_n}$ is a chordless cycle in the Gaifman graph of $\mathfrak{A} \otimes \mathbb{G}$. W.l.o.g. the representatives (g_i, s_i, a_i) are chosen such that $[g_{i+1}, s_{i+1}]$ is a hyperedge linking $[g_i, s_i, a_i]$ and $[g_{i+1}, s_{i+1}, a_{i+1}]$ in $\mathfrak{A} \otimes \mathbb{G}$. I.e., there is a path w from s_i to s_{i+1} in I consisting of edges from $\alpha_i := E_{a_i}$ such that $g_{i+1} = g_i \cdot w^{\mathbb{G}}$. In particular, $g_i^{-1} \cdot g_{i+1} \in \mathbb{G}_{\alpha_i}$. We claim that $(g_i)_{i \in \mathbb{Z}_n}$ with the colouring $(\alpha_i)_{i \in \mathbb{Z}_n}$ is a coloured cycle in \mathbb{G} in the sense of Definition 2.12. If so, $n > N$ follows, since \mathbb{G} is N -acyclic.

In connection with $(g_i)_{i \in \mathbb{Z}_n}$ and $(\alpha_i)_{i \in \mathbb{Z}_n}$ it essentially just remains to check the colouring condition

$$g_i \mathbb{G}_{\alpha_i \alpha_{i-1}} \cap g_{i+1} \mathbb{G}_{\alpha_i \alpha_{i+1}} = \emptyset.$$

Suppose, for contradiction, that there is some $k \in g_i \mathbb{G}_{\alpha_i \alpha_{i-1}} \cap g_{i+1} \mathbb{G}_{\alpha_i \alpha_{i+1}}$, and let $t \in S$ be such that $k \in G_{*t}$. We show that this situation implies that $[g_{i-1}, s_{i-1}, a_{i-1}]$ and $[g_{i+1}, s_{i+1}, a_{i+1}]$ are linked by a chord induced by the hyperedge $[t, k]$:

(a) Since $k \in g_i \mathbb{G}_{\alpha_i \alpha_{i-1}}$, there is some path w_1 from s_i to t consisting of edges in $\alpha_i \cap \alpha_{i-1}$ such that $k = g_i \cdot w_1^{\mathbb{G}}$; as there also is a path w from s_{i-1} to s_i in I consisting of edges from α_{i-1} such that $g_i = g_{i-1} \cdot w^{\mathbb{G}}$, it follows that there is a path w_2 from s_{i-1} to t consisting of edges in α_{i-1} such that $k = g_{i-1} \cdot w_2^{\mathbb{G}}$. So $[g_{i-1}, s_{i-1}, a_{i-1}] \in [t, k]$.

(b) Since $k \in g_{i+1} \mathbb{G}_{\alpha_i \alpha_{i+1}}$, there is some path w_3 from s_{i+1} to t consisting of edges in $\alpha_i \cap \alpha_{i+1} \subseteq \alpha_{i+1}$ such that $k = g_{i+1} \cdot w_3^{\mathbb{G}}$; so $[g_{i+1}, s_{i+1}, a_{i+1}] \in [t, k]$.

(a) and (b) together imply that the given cycle was not chordless after all. \square

3.2.2 Conformality in reduced products

Lemma 3.8. *Let \mathfrak{A} be a hypergraph with intersection graph $I(\mathfrak{A}) =: I$, \mathbb{G} an N -acyclic I -groupoid. Then $\mathfrak{A} \otimes \mathbb{G}$ is N -conformal.*

Proof. Suppose that $X := \{[g_i, s_i, a_i] : i \in n\}$ is a clique of the Gaifman graph of $\mathfrak{A} \otimes \mathbb{G}$ that is not contained in a hyperedge of $\mathfrak{A} \otimes \mathbb{G}$, but such that every subset of $n-1$ vertices is contained in a hyperedge of $\mathfrak{A} \otimes \mathbb{G}$: for $i \in n$, choose a hyperedge $[k_i, t_i]$ such that $X_i := \{[g_j, s_j, a_j] : j \neq i\} \subseteq [k_i, t_i]$. Let $h_i := k_i^{-1} \cdot k_{i+1}$ and $\alpha_i := \bigcap_{j \neq i, i+1} E_{a_j}$. Note that $k_i^{-1} \cdot k_{i+1} \in \mathbb{G}_{\alpha_i}$. We claim that $(k_i)_{i \in \mathbb{Z}_n}$ with the colouring $(\alpha_i)_{i \in \mathbb{Z}_n}$ is a coloured cycle in \mathbb{G} in the sense of Definition 2.12. It follows that $n > N$, as desired.

Suppose, for contradiction, that $k \in k_i \mathbb{G}_{\alpha_i \alpha_{i-1}} \cap k_{i+1} \mathbb{G}_{\alpha_i \alpha_{i+1}}$ for some i . Let $t \in S$ be such that $k \in G_{*t}$. We show that $X \subseteq [k, t]$ would follow.

Since $k \in k_i \mathbb{G}_{\alpha_i \alpha_{i-1}}$, and $[g_j, s_j, a_j] \in [k_i, t_i]$, i.e., $[g_j, s_j, a_j] = [k_i, t_i, a_j]$, for all $j \neq i$, clearly $[g_j, s_j, a_j] \in [k, t]$ for $j \neq i$ (note that $\alpha_i \cap \alpha_{i-1} = \bigcap_{j \neq i} E_{a_j}$).

It therefore remains to argue that also $[g_i, s_i, a_i] \in [k, t]$. Note that $k \in k_{i+1} \mathbb{G}_{\alpha_i \alpha_{i+1}}$ and that $\alpha_i \cap \alpha_{i+1} = \bigcap_{j \neq i+1} E_{a_j}$. In particular, generators in

$\alpha_i \cap \alpha_{i+1}$ preserve a_i . Since $[g_i, s_i, a_i] \in [k_{i+1}, t_{i+1}]$, we have that $[g_i, s_i, a_i] = [k_{i+1}, t_i, a_i]$, and thus $[g_i, s_i, a_i] \in [k, t]$ follows from the fact that $k_{i+1}^{-1} \cdot k \in \mathbb{G}_{\alpha_i \alpha_{i+1}}$. \square

3.3 N-acyclic covers

Combining the above, we obtain the following by application of the reduced product construction $\mathfrak{A} \otimes \mathbb{G}$ for suitably acyclic groupoids \mathbb{G} .

Proposition 3.9. *For every $N \in \mathbb{N}$, every finite hypergraph admits a cover by a finite hypergraph that is N -acyclic.*

3.4 Variations on unfoldings

3.4.1 Unfoldings w.r.t. link structure

We explore the analogue of the construction of hypergraph covers in the situation where hyperedge overlaps in the cover may be limited to specified subsets of the overlaps realised in the given hypergraph. This variation is of interest in contexts where even the overlap between two hyperedges s and s' may give rise to problematic ‘cycles’ (w.r.t. some underlying relational content say), so that one may want to use two different copies of s' that overlap with s in different subsets of $s \cap s'$ in order to break such a ‘cycle’. Of course any such variation produces not a proper cover, since a cover must reproduce overlaps in full. We therefore look at modifications of the *back*-property; and at specifications of the required overlaps given in terms of a link multi-graph I rather than the intersection graph used so far.

Definition 3.10. A *link structure* for a hypergraph $\mathfrak{A} = (A, S)$ is a directed loop-free multi-graph $I = (S, E)$ with converse edges, where $E = (E[s, s'])_{s \neq s'}$ is such that every edge $e \in E[s, s']$ and its converse $e^{-1} \in E[s', s]$ are associated with the same non-empty subset $d[e] = d[e^{-1}] \subseteq s \cap s'$.

Intuitively, the set $d[e] \subseteq s \cap s'$ for $e \in E[s, s']$ is the domain in which copies of s and s' are meant to overlap according to e . The intended covering task is captured by the following definition, which characteristically modifies the back-requirement of Definition 3.3.

Definition 3.11. Let $\mathfrak{A} = (A, S)$ and $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ be hypergraphs, $I = (S, E)$ a link structure for \mathfrak{A} . A hypergraph homomorphism $h: \hat{\mathfrak{A}} \rightarrow \mathfrak{A}$ is an *I -cover* if it satisfies the back-property w.r.t. links specified in I : for every $h(\hat{s}) = s \in S$, $s' \in S$ and $e \in E[s, s']$ in I , there is some $\hat{s}' \in \hat{S}$ such that $h(\hat{s}') = s'$ and $h(\hat{s} \cap \hat{s}') = d[e] \subseteq s \cap s'$.

We note that these definitions reproduce the original definition of a hypergraph cover as the special case of link structure $I = I(\mathfrak{A})$ with $d[(s, s')] = s \cap s'$.

We follow the same recipe as before, and produce I -covers as reduced products with I -groupoids. For a hypergraph \mathfrak{A} , link structure I and an I -groupoid

\mathbb{G} , we define a hypergraph

$$\mathfrak{A} \otimes_E \mathbb{G} = (\hat{A}, \hat{S})$$

as a reduced product as follows. The vertex set \hat{A} of $\mathfrak{A} \otimes_E \mathbb{G}$ is the quotient of the disjoint union of hyperedges s of \mathfrak{A} tagged by groupoid elements $g \in G_{*s}$, $\bigcup_{s \in S, g \in G_{*s}} \{g\} \times \{s\} \times s$ w.r.t. the equivalence relation induced by identifications

$$(g, s, a) \approx (ge, s', a) \quad \text{for } e \in E(s, s') \text{ with } a \in d[e].$$

So now (g_1, s_1, a) is identified with (g_2, s_2, a) if, and only if, $g_2 = g_1 \cdot w^{\mathbb{G}}$ for some path w from s_1 to s_2 in I consisting of edges e for which $a \in d[e]$. Denoting the equivalence class of a triple (g, s, a) as $[g, s, a]$ as before, the hyperedges of $\mathfrak{A} \otimes_E \mathbb{G} = (\hat{A}, \hat{S})$ are given by

$$\hat{S} = \{[g, s] : s \in S, g \in G_{*s}\} \text{ where } [g, s] := \{[g, s, a] : a \in s\} \subseteq \hat{A}.$$

The cover homomorphism π is the natural projection

$$\pi : [g, s, a] \mapsto a,$$

and it is easy to see that $\pi : \hat{\mathfrak{A}} \rightarrow \mathfrak{A}$ satisfies the back-property for overlaps specified in I , whence it is an I -cover.

Lemma 3.12. *For an I -groupoid \mathbb{G} that is 2-acyclic:*

- (a) *the overlap between any two hyperedges $[g_0, s]$ and $[g_1, t]$ of $\mathfrak{A} \otimes_E \mathbb{G}$ is the pre-image of the set $\bigcap_i d[e_i]$ for some $w = e_1 \dots e_n \in E_{st}^*$ for which $g_1 = g_0 \cdot w^{\mathbb{G}}$. More specifically,*

$$g_0^{-1} \cdot g_1 \in \mathbb{G}_\alpha \quad \text{where } \alpha = \{e \in E : \pi([g_0, s] \cap [g_1, t]) \subseteq d[e]\}.$$

- (b) *For $e \in E[s, s']$, the overlap between hyperedges $[g, s]$ and $[g \cdot e, s']$ is precisely the subset*

$$[g, d[e]] = \{[g, a] : a \in d[e]\} = \{[g \cdot e, a] : a \in d[e]\} = [g \cdot e, d[e]].$$

Proof. (a) For $a \in A$, let $E_a := \{e \in E : a \in d[e]\} \subseteq E$ be the subset of those edges that preserve a , and similarly $E_d := \{e \in E : d \subseteq d[e]\}$ for any subset $d \subseteq A$. As before, \mathbb{G}_α denotes the sub-groupoid generated by $\alpha = \alpha^{-1} \subseteq E$. For the purposes of this proof, let $\mathbb{G}_d := \mathbb{G}_{E_d}$ and $\mathbb{G}_a := \mathbb{G}_{E_a}$. Note that \mathbb{G}_a is the sub-groupoid whose cosets determine the spread of hyperedges of $\mathfrak{A} \otimes_E \mathbb{G}$ that share the same element above a .

Let $h := g_0^{-1} \cdot g_1$ and $d := \pi([g_0, s] \cap [g_1, t])$. We want to show that $h \in \mathbb{G}_d$. By Observation 2.18, and since \mathbb{G} is 2-acyclic, $\mathbb{G}_d = \bigcap_{a \in d} \mathbb{G}_a$, and it is clear that $h \in \mathbb{G}_a$ for each $a \in d$, since $[g_0, s, a] \in [g_0, s] \cap [g_1, t]$.

For (b) suppose the intersection $[g, s] \cap [g \cdot e, s']$ properly contained the subset corresponding to $d[e]$. Then, by (a), there must some $w \in E_{ss'}^*$ s.t. $w^{\mathbb{G}} = g^{-1} \cdot (g \cdot e) = e$ but $d[e] \subsetneq \bigcap d[e_i]$. It follows that $e \in \mathbb{G}_\alpha$ for $\alpha = E \setminus \{e, e^{-1}\}$ as well as $e \in \mathbb{G}_{\{e, e^{-1}\}}$, which, for 2-acyclic \mathbb{G} , contradicts Observation 2.18. \square

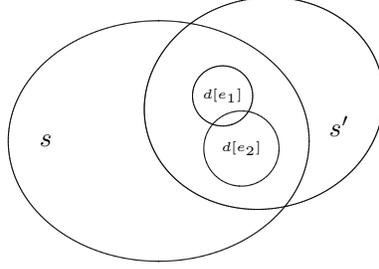


Figure 4: Link structure.

It is not hard to see, by arguments strictly analogous to those given for covers in Section 3.1, that $\mathfrak{A} \otimes_E \mathbb{G}$ is N -acyclic for N -acyclic \mathbb{G} .

Proposition 3.13. *Any finite hypergraph \mathfrak{A} with specified link structure I admits, for any $N \in \mathbb{N}$, an I -cover by a finite N -acyclic hypergraph $\hat{\mathfrak{A}}$.*

3.4.2 Unfoldings of I-graphs

Rather than reduced products with a given hypergraph, we now define a hypergraph $H \otimes \mathbb{G}$ as a reduced product based on an I -graph H and I -groupoid \mathbb{G} .

Let $H = (V, (V_s), (R_e))$ be a not necessarily complete I -graph, $I = (S, E)$ as before. For $w \in E_{s_t}^*$ we look at the composition ρ_w^H of the partial maps ρ_e^H induced by the partial matchings R_e of H as partial bijections from V_s to V_t (admitting also the empty partial map).

For an I -groupoid \mathbb{G} we define a hypergraph $\mathfrak{B} := (B, \hat{S})$ as a reduced product of H and \mathbb{G} as follows. The vertex set B is the quotient of the disjoint union of the sets V_s in H , tagged by groupoid elements $g \in G_{*s}$ and (redundantly) by s itself,

$$\bigcup_{s \in S, g \in G_{*s}} \{g\} \times \{s\} \times V_s$$

w.r.t. the equivalence relation \approx induced by

$$(g, s, u) \approx (g', s', v) \quad \text{if } g \in G_{*s}, e \in E[s, s'] \text{ and } (u, v) \in R_e.$$

The identifications induced by \approx are as follows. For $g_i \in G_{*s_i}$ and $v_i \in V_{s_i}$, (g_1, s_1, v_1) is identified with (g_2, s_2, v_2) in the quotient if, and only if, there is a path $w \in E_{s_1 s_2}^*$ such that $\rho_w^H(v_1) = v_2$ in H and $g_2 = g_1 \cdot w^{\mathbb{G}}$ in \mathbb{G} . Denoting the equivalence class of (g, s, v) as $[g, s, v]$, we define the vertex set B of $\mathfrak{B} = H \otimes \mathbb{G} = (B, \hat{S})$ as

$$B := \{[g, s, v]: s \in S, g \in G_{*s}, v \in V_s\},$$

and the hyperedges of \mathfrak{B} to be the subsets represented by the natural copies of the patches V_s of H :

$$\hat{S} = \{[g, s]: g \in \mathbb{G}\} \quad \text{where, for } g \in G_{*s}, \quad [g, s] := \{[g, s, v]: v \in V_s\} \subseteq B.$$

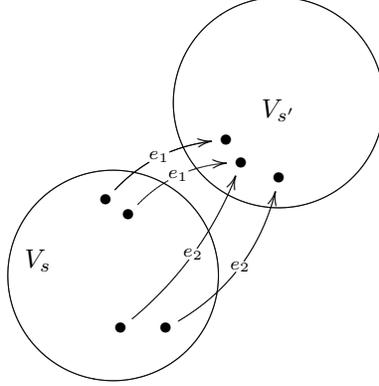


Figure 5: Overlap pattern.

Lemma 3.14. *If \mathbb{G} is compatible with the I -graph $H = (V, (V_s), (E_s))$, then the natural projection $\pi_{[g,s]}: [g, v] \mapsto v$ is well-defined in restriction to each hyperedge $[g, s]$ of $H \otimes \mathbb{G}$, and relates the hyperedge $[g, s] = \{[g, s, v]: v \in V_s\}$ bijectively to V_s .*

Proof. It suffices to show that $[g, s, v] = [g, s, v']$ implies $v = v'$, which shows that $\pi_{[g,s]}$ is well-defined. By compatibility of \mathbb{G} with H , we find that for $w \in E_{ss}^*$, $w^{\mathbb{G}} = 1$ implies $\rho_w^H \subseteq \text{id}_{V_s}$. This follows from the observation that, where the composition ρ_w^H is defined, it agrees with the operation of ρ_w on the completion \bar{H} , which gives rise to $\text{cym}(H)$. \square

3.4.3 Realisations of overlap patterns

Consider an I -graph $H = (V, (V_s), (R_e))$ as a specification of overlaps to be realised between isomorphic copies of the sets V_s according to identifications induced by the partial matchings R_e or $\rho_e = \rho_e^H$. Reduced products $H \otimes \mathbb{G}$ with I -groupoids \mathbb{G} that are compatible with H are a first approximation: according to Lemma 3.14, the hyperedges of $H \otimes \mathbb{G}$ individually project bijectively onto the respective sets V_s . It is not clear, however, that hyperedges $[g_1, s]$ and $[g_2, t]$, related to V_s and V_t , do not overlap to a greater extent than specified in H . We formalise the desired congruence with the overlap pattern specified in H as follows.

Definition 3.15. Let $I = (S, E)$ be a directed loop-free multi-graph, $H = (V, (V_s), (R_e))$ an I -graph with induced partial bijections ρ_w^H between V_s and V_t for $w \in E_{st}^*$. A hypergraph $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ is a *realisation* of the overlap pattern specified by H , if there is a map $\pi: \hat{S} \rightarrow S$ and a matching family of bijections

$$\pi_{\hat{s}}: \hat{s} \longrightarrow V_s, \text{ for } \hat{s} \in \hat{S} \text{ with } \pi(\hat{s}) = s,$$

such that for all $\hat{s}, \hat{t} \in \hat{S}$ s.t. $\pi(\hat{s}) = s$, $\pi(\hat{t}) = t$, and for every $e \in E[s, s']$:

- (i) if $\hat{s} \cap \hat{t} \neq \emptyset$, then $\pi_{\hat{t}} \circ \pi_{\hat{s}}^{-1} = \rho_w^H$ for some $w \in E_{st}^*$;
- (ii) there is some \hat{s}' such that $\pi(\hat{s}') = s'$ and $\rho_e^H = \pi_{\hat{s}'} \circ \pi_{\hat{s}}^{-1}$.

$$\begin{array}{ccc}
& \hat{s} \cap \hat{t} & \\
\pi_{\hat{s}} \swarrow & & \searrow \pi_{\hat{t}} \\
V_s & \xrightarrow{\rho_w^H} & V_t
\end{array}$$

Some comment on the definition: condition (i) says that all overlaps between hyperedges realised in $\hat{\mathfrak{A}}$ are induced by overlaps specified in H – allowing for compositions along chains of overlaps; condition (ii) says that all those local overlaps that should be realised according to H are indeed realised at corresponding sites in $\hat{\mathfrak{A}}$.

By Lemma 3.14, $\tilde{\mathfrak{A}} = (A, \tilde{S}) := H \otimes \mathbb{G}$ for an I -groupoid \mathbb{G} that is compatible with H , has bijective projections $\pi_{\tilde{s}}: \tilde{s} = [g, s] \rightarrow V_s$ for $g \in G_{*s}$. In order to realise conditions (i) and (ii), we want to apply an unfolding of the form

$$\hat{\mathfrak{A}} = \tilde{\mathfrak{A}} \otimes_{\tilde{E}} \tilde{\mathbb{G}}$$

w.r.t. a suitable link structure $\tilde{I} = (\tilde{S}, \tilde{E})$ and a 2-acyclic \tilde{I} -groupoid $\tilde{\mathbb{G}}$.

With $e \in E[s, s']$ (from the original multi-graph $I = (S, E)$) we associate edges $\tilde{e} \in \tilde{E}[\tilde{s}, \tilde{s}']$ from $\tilde{s} = [g, s] \in \tilde{S}$ to $\tilde{s}' = [g \cdot e, s'] \in \tilde{S}$. These hyperedges \tilde{s} and \tilde{s}' have bijective projections $\pi_{\tilde{s}}: [g, s] \rightarrow V_s$ and $\pi_{\tilde{s}'}: [g \cdot e, s'] \rightarrow V_{s'}$.⁵ In order to enforce conditions (i) and (ii) we specify $\tilde{e} \in \tilde{E}[\tilde{s}, \tilde{s}']$ through

$$d[\tilde{e}] := \{[g, s, v]: v \in \text{dom}(\rho_e^H)\} = \{[g \cdot e, s', \rho_e^H(v)]: v \in \text{dom}(\rho_e^H)\} \subseteq \tilde{s} \cap \tilde{s}'.$$

A hyperedge of $\hat{\mathfrak{A}}$ is of the form $\hat{s} = [\tilde{g}, \tilde{s}]$ where $\tilde{g} \in \tilde{G}_{*\tilde{s}}$. From $\hat{\mathfrak{A}} = \tilde{\mathfrak{A}} \otimes_{\tilde{E}} \tilde{\mathbb{G}}$, this hyperedge \hat{s} bijectively projects onto the hyperedge \tilde{s} of $\tilde{\mathfrak{A}} = H \otimes \mathbb{G}$, and from there bijectively onto the appropriate V_s .

Claim 3.16. *If $\tilde{\mathbb{G}}$ is a 2-acyclic \tilde{I} -groupoid, then $\hat{\mathfrak{A}} = \tilde{\mathfrak{A}} \otimes_{\tilde{E}} \tilde{\mathbb{G}}$, with the natural projections as outlined, is a realisation of the overlap pattern specified by H .*

Proof. It remains to argue that conditions (i) and (ii) of the definition are satisfied. By construction of the reduced product $\tilde{\mathfrak{A}} \otimes_{\tilde{E}} \tilde{\mathbb{G}}$ it is clear that, in the situation of (ii), the \tilde{e} -edge at \tilde{s} in \tilde{I} produces a hyperedge \hat{s}' that overlaps with \hat{s} as prescribed by the overlap $d[\tilde{e}]$ between \tilde{s} and \tilde{s}' . That the overlap is no larger than specified in $d[\tilde{e}]$, follows from part (b) of Lemma 3.12. For (i), we use part (a) of Lemma 3.12. \square

Proposition 3.17. *For every directed loop-free multi-graph I and I -graph H , there is a finite hypergraph $\hat{\mathfrak{A}}$ that realises the overlap pattern specified by H . For a given threshold $N \in \mathbb{N}$, such realisations $\hat{\mathfrak{A}}$ can be chosen N -acyclic.*

⁵Formally we get a whole family of edges \tilde{e} , one for each $g \in G_{*s}$, but do not distinguish them by distinct names, because the relevant loci are always specified implicitly when we speak of $\tilde{e} \in \tilde{E}[[g], [g \cdot e]]$.

4 Symmetries

4.1 Groupoidal symmetries

Consider a hypergraph cover $\hat{\mathfrak{A}} = \mathfrak{A} \otimes \mathbb{G}$ of the hypergraph \mathfrak{A} by means of a reduced product with an I -groupoid \mathbb{G} , for $I = I(\mathfrak{A})$. Any such cover has characteristic symmetries within its fibres of hyperedges. Symmetries are described in terms of automorphisms.

Definition 4.1. An *automorphism* of a hypergraph $\mathfrak{A} = (A, S)$ is a bijection $\eta: A \rightarrow A$ such that for every $s \subseteq A$: $\eta(s) := \{\eta(a) : a \in s\} \in S$ if, and only if, $s \in S$. We also denote the induced bijection on S as η .

Lemma 4.2. Let $\pi: \mathfrak{A} \otimes \mathbb{G} \rightarrow \mathfrak{A}$ be the natural cover of a hypergraph $\mathfrak{A} = (A, S)$ by a reduced product with some I -groupoid \mathbb{G} where $I = I(\mathfrak{A})$. Then any two pre-images of the same hyperedge $s \in S$ are related by an automorphism of $\hat{\mathfrak{A}} = \mathfrak{A} \otimes \mathbb{G}$ that commutes with π .

Proof. W.l.o.g. I is connected; otherwise \mathfrak{A} and \mathbb{G} consist of disjoint components, which can be treated separately.

Let $[g_i, s]$ for $g_i \in G_{*s}$, $i = 1, 2$ be any two hyperedges of $\hat{\mathfrak{A}}$ above the same $s \in S$. Let $g := g_2 \cdot g_1^{-1} \in G_{ss}$. For $t \in S \setminus \{s\}$ choose $g_t \in G_{st}$ and put $g_s := 1_s$. Then the following is an automorphism of the Cayley graph of \mathbb{G} :

$$\begin{aligned} \eta: G &\longrightarrow G \\ k \in G_{tt'} &\longmapsto g_t^{-1} \cdot g \cdot g_t \cdot k \in G_{tt'}, \end{aligned}$$

which maps g_1 to g_2 . It follows that the natural extension of this map to $\hat{\mathfrak{A}}$,

$$\begin{aligned} \hat{\eta}: \mathfrak{A} \otimes \mathbb{G} &\longrightarrow \mathfrak{A} \otimes \mathbb{G} \\ [g, s, a] &\longmapsto [\eta(g), s, a], \end{aligned}$$

is well-defined and an automorphism of $\hat{\mathfrak{A}}$ as desired, since it maps $[g_1, s]$ to $[g_2, s]$ and is compatible with π . \square

The arguments presented similarly support the case of reduced products $\mathfrak{A} \otimes_E \mathbb{G}$ of a hypergraph with specified link structure, and the case of reduced products $H \otimes \mathbb{G}$ between I -graphs and I -groupoids, as discussed in Section 3.4.

Corollary 4.3. The reduced products $\mathfrak{A} \otimes_E \mathbb{G}$ and $H \otimes \mathbb{G}$, where \mathbb{G} is compatible with H , admit automorphisms that relate any given two hyperedges above the same hyperedge s of \mathfrak{A} , or above the same V_s of H , and are compatible with the projection $\pi: \mathfrak{A} \otimes_E \mathbb{G} \rightarrow \mathfrak{A}$, or with the local projections $\pi_{\hat{s}}: \hat{s} \rightarrow V_s$, respectively.

4.2 Lifting structural symmetries

Besides the vertical symmetries within fibres there is an obvious concern relating to the compatibility of reduced products with automorphisms of the given structure. Our aim is to show that the constructions are sufficiently natural

or canonical to allow us to lift symmetries of a hypergraph \mathfrak{A} to its covers by $\mathfrak{A} \otimes \mathbb{G}$, of a hypergraph \mathfrak{A} with specified link structure to its unfoldings $\mathfrak{A} \otimes_E \mathbb{G}$, and of an I -graph H to its unfolding $H \otimes \mathbb{G}$.

For this we need to look at symmetries of the associated $I = (S, E)$, which however are not automorphisms of the (multi-)graph I if we think of the edge colours as fixed as we have so far. The relevant automorphisms of \mathfrak{A} or H relate to certain automorphisms of the multi-graphs $I = (S, E)$ which may permute vertices and edge and vertex colours.

The following definitions are just the natural ones. Because of the different levels involved in symmetries and automorphisms it may make sense to state them explicitly.

Definition 4.4. Let $I = (S, E)$ be any loop-free directed multi-graph with converse edges $e^{-1} \in E$ for every $e \in E$.

A *symmetry of I* is an automorphism of the associated two-sorted incidence structure, i.e., a pair $\eta = \eta^I = (\eta^S, \eta^E)$ of bijections $\eta^S: S \rightarrow S$ and $\eta^E: E \rightarrow E$, such that $\eta^E(e) \in E[\eta^S(s), \eta^S(t)]$ iff $e \in E[s, t]$.⁶

A symmetry $\eta^I = (\eta^S, \eta^E)$ of I is a *symmetry of the link structure $I = (S, E)$* of a hypergraph $\mathfrak{A} = (A, S)$ if it is induced by a hypergraph automorphism $\eta^{\mathfrak{A}}: \mathfrak{A} \rightarrow \mathfrak{A}$ through the natural stipulations for $\eta^S: S \rightarrow S$ (operation of $\eta^{\mathfrak{A}}$ on subsets of its domain A), $\eta^E: E \rightarrow E$ (mapping $e \in E[s, s']$ to $\eta(e) \in E[\eta(s), \eta(s')]$ where $d[\eta(e)] = \eta(d[e])$ through operation of η on subsets $d \subseteq A$). In this scenario we think of η as the triple $\eta = (\eta^{\mathfrak{A}}, \eta^S, \eta^E)$.

A symmetry $\eta = (\eta^S, \eta^E)$ of I is a *symmetry of an I -graph $H = (V, (V_s), (R_e))$* if it is induced by a bijection $\eta^H: V \rightarrow V$ such that for all $s \in S$ and $e \in E$:

- (i) $\eta^H(V_s) := \{\eta^H(v) : v \in V_s\} = V_{\eta^S(s)}$, and
- (ii) $\eta^H(R_e) := \{(\eta^H(v), \eta^H(v')) : (v, v') \in R_e\} = R_{\eta^E(e)}$.

In this scenario we think of η as the triple $\eta = (\eta^H, \eta^S, \eta^E)$.

Note that any symmetry of I is compatible with converses: $\eta(e^{-1}) = (\eta(e))^{-1}$.

Definition 4.5. A *symmetry of an I -groupoid \mathbb{G}* with generators $e \in E$ is a bijection $\eta^{\mathbb{G}}: G \rightarrow G$ with an induced symmetry $\eta^I = (\eta^S, \eta^E)$ of I , such that for all $e \in E[s, s']$, $\eta^{\mathbb{G}}$ maps the generator $e \in G_{ss'}$ to the generator $\eta(e) \in G_{\eta(s)\eta(s')}$ and, correspondingly, for all $s \in S$ and $g_1 \in G_{st}, g_2 \in G_{tu}$:

- (i) $\eta^{\mathbb{G}}(1_s) = 1_{\eta(s)}$;
- (ii) $\eta^{\mathbb{G}}(g_1 \cdot g_2) = \eta^{\mathbb{G}}(g_1) \cdot \eta^{\mathbb{G}}(g_2)$.

In this scenario we think of η as $\eta = (\eta^{\mathbb{G}}, \eta^I)$.

It follows from the definition that $\eta^{\mathbb{G}}(g) \in G_{\eta(s)\eta(t)}$ for all $g \in G_{st}$ and that, for $w = e_1 \dots e_n \in E_{st}^*$, $\eta^{\mathbb{G}}(w^{\mathbb{G}}) = (\eta(w))^{\mathbb{G}}$ where $\eta(w) = \eta(e_1) \dots \eta(e_n) \in E_{\eta(s)\eta(t)}^*$.

⁶We often write just η to denote the different incarnations of η , and use superscripts to highlight different domains where necessary.

It is obvious that the Cayley graph of an I -groupoid \mathbb{G} with symmetry η has η as a symmetry of the I -graph. Conversely, any symmetry η of an I -graph H lifts to an automorphism of the I -groupoid $\mathbb{G} = \text{cym}(H)$.

An automorphism $\eta^{\mathfrak{A}}$ of a hypergraph \mathfrak{A} induces a symmetry $\eta^I = (\eta^S, \eta^E)$ of $I = I(\mathfrak{A})$. If the I -groupoid \mathbb{G} has a symmetry $\eta = (\eta^{\mathbb{G}}, \eta^I)$ with the same underlying symmetry η^I , then the cover $\pi: \mathfrak{A} \otimes \mathbb{G} \rightarrow \mathfrak{A}$ carries a corresponding symmetry that is both an automorphism of the covering hypergraph $\hat{\mathfrak{A}} = \mathfrak{A} \otimes \mathbb{G}$ and compatible with the given automorphism of \mathfrak{A} via π , in the sense that the following diagram commutes.

$$\begin{array}{ccc} \mathfrak{A} \otimes \mathbb{G} & \xrightarrow{\eta^{\hat{\mathfrak{A}}}} & \mathfrak{A} \otimes \mathbb{G} \\ \downarrow \pi & & \downarrow \pi \\ \mathfrak{A} & \xrightarrow{\eta^{\mathfrak{A}}} & \mathfrak{A} \end{array}$$

At the level of $\mathfrak{A} \otimes \mathbb{G}$, the automorphism η operates according to

$$\eta^{\mathfrak{A} \otimes \mathbb{G}}: [g, s, a] \mapsto [\eta^{\mathbb{G}}(g), \eta^S(s), \eta^{\mathfrak{A}}(a)],$$

where all the different incarnations of η are highlighted by superscripts.

Upon re-inspection, the construction steps in the generation of N -acyclic I -groupoids are naturally compatible with any symmetry of the given I . The induction underlying Proposition 2.17 is based on the number of generators in the sub-groupoids, with an individual induction step according to Lemma 2.16 on the amalgamation of chains of sub-groupoids. All these notions are entirely symmetric w.r.t. any symmetries of I .

Corollary 4.6. *If $I = (S, E)$ is a finite loop-free directed multi-graph, then there is, for every $N \in \mathbb{N}$, some finite N -acyclic I -groupoid \mathbb{G} whose symmetries induce all the symmetries of I (i.e., all symmetries of I can be lifted to \mathbb{G}).*

From Proposition 3.9 we thus further obtain the following.

Corollary 4.7. *Any finite hypergraph \mathfrak{A} admits, for $N \in \mathbb{N}$, finite N -acyclic covers by reduced products with finite N -acyclic I -groupoids, $\pi: \hat{\mathfrak{A}} = \mathfrak{A} \otimes \mathbb{G} \rightarrow \mathfrak{A}$ that are compatible with the automorphism group of \mathfrak{A} in the sense that every automorphism $\eta^{\mathfrak{A}}$ of \mathfrak{A} lifts to an automorphism $\eta^{\hat{\mathfrak{A}}}$ such that $\pi \circ \eta^{\hat{\mathfrak{A}}} = \eta^{\mathfrak{A}} \circ \pi$.*

It is clear from the arguments presented that reduced products $\mathfrak{A} \otimes_E \mathbb{G}$ that unfold a hypergraph \mathfrak{A} w.r.t. a specified link structure and reduced products $H \otimes \mathbb{G}$ that unfold I -graphs H , as discussed in Section 3.4, support the same kinds of symmetries as $\mathfrak{A} \otimes \mathbb{G}$.

Corollary 4.8. *Every symmetry η of \mathfrak{A} , the link structure I and \mathbb{G} gives rise to an automorphism of $\mathfrak{A} \otimes_E \mathbb{G}$. Every symmetry η of H , I and \mathbb{G} gives rise to an automorphism of $H \otimes \mathbb{G}$.*

Corollary 4.9. *For any finite loop-free multi-graph $I = (S, E)$ and I -graph $H = (V, (V_s), (R_e))$, realisations $\hat{\mathfrak{A}}$ as obtained in Proposition 3.17 can be chosen so that all symmetries of H lift to automorphisms of $\hat{\mathfrak{A}}$. Moreover, for any two hyperedges \hat{s}_1 and \hat{s}_2 of $\hat{\mathfrak{A}}$ that bijectively project to the same V_s for $s = \pi(\hat{s}_1) = \pi(\hat{s}_2)$ via $\pi_{\hat{s}_i}: \hat{s}_i \rightarrow V_s$, there is a ‘vertical’ automorphism of $\hat{\mathfrak{A}}$ that is compatible with these projections.*

Proof. It suffices to choose the two groupoids in the construction of realisations according to Section 3.4.3 sufficiently symmetric. Recall the two-level construction

$$H \longrightarrow \tilde{\mathfrak{A}} = H \otimes \mathbb{G} \longrightarrow \hat{\mathfrak{A}} = \tilde{\mathfrak{A}} \otimes_{\tilde{E}} \tilde{\mathbb{G}},$$

where \mathbb{G} is an I -groupoid that is compatible with H , and $\tilde{\mathbb{G}}$ a 2-acyclic \tilde{I} -groupoid based on the link structure $\tilde{I} = (\tilde{S}, \tilde{E})$ induced by edges $e \in E[s, s']$ between hyperedges $\tilde{s} = [g, s]$ and $\tilde{s}' = [g \cdot e, s']$ with overlaps $d[\tilde{e}]$ given by ρ_e^H .

If \mathbb{G} is chosen symmetric w.r.t. H and I , it lifts all symmetries of H to automorphisms of $\tilde{\mathfrak{A}}$, and also $\tilde{\mathfrak{A}}$ has groupoidal symmetries that relate any two hyperedges above the same V_s in a manner that is compatible with their projections to V_s . These automorphisms of $\tilde{\mathfrak{A}}$ give rise to symmetries of \tilde{I} , and if $\tilde{\mathbb{G}}$ is chosen symmetric w.r.t. all automorphisms of $\tilde{\mathfrak{A}}$ and all (induced) symmetries \tilde{I} , then $\tilde{\mathfrak{A}}$ will lift all automorphisms of $\tilde{\mathfrak{A}}$, and hence in particular all automorphisms of H .

Two hyperedges $\hat{s}_1 = [\tilde{g}_1, \tilde{s}_1]$ and $\hat{s}_2 = [\tilde{g}_2, \tilde{s}_2]$ above the same $\tilde{s} = \tilde{s}_1 = \tilde{s}_2$ are clearly related by a groupoidal ‘vertical’ automorphism of $\tilde{\mathfrak{A}} = \tilde{\mathfrak{A}} \otimes_{\tilde{E}} \tilde{\mathbb{G}}$, that is compatible with the projection to \tilde{s} and hence to V_s .

Two hyperedges $\hat{s}_1 = [\tilde{g}_1, \tilde{s}_1]$ and $\hat{s}_2 = [\tilde{g}_2, \tilde{s}_2]$ of $\hat{\mathfrak{A}}$ that project onto the same V_s via $\pi_{\hat{s}_i}: \hat{s}_i \rightarrow V_s$ but for distinct $\tilde{s}_1 \neq \tilde{s}_2$, we look at the projection $\pi: \hat{\mathfrak{A}} = \tilde{\mathfrak{A}} \otimes_{\tilde{E}} \tilde{\mathbb{G}} \rightarrow \tilde{\mathfrak{A}}$ to find that their images \tilde{s}_1 and \tilde{s}_2 are related by an automorphism of $\tilde{\mathfrak{A}}$. That automorphism induces a symmetry of \tilde{I} and lifts to an automorphism of $\hat{\mathfrak{A}}$ that maps $\hat{s}_1 = [\tilde{g}_1, \tilde{s}_1]$ to some $\hat{s}'_2 = [\tilde{g}'_2, \tilde{s}_2]$ above \tilde{s}_2 . But then $\hat{\mathfrak{A}} = \tilde{\mathfrak{A}} \otimes_{\tilde{E}} \tilde{\mathbb{G}}$ also has a ‘vertical’ groupoidal automorphism that maps $\hat{s}'_2 = [\tilde{g}'_2, \tilde{s}_2]$ to $\hat{s}_2 = [\tilde{g}_2, \tilde{s}_2]$ (cf. Corollary 4.3). The composition of these two automorphisms therefore maps $\hat{s}_1 = [\tilde{g}_1, \tilde{s}_1]$ to $\hat{s}_2 = [\tilde{g}_2, \tilde{s}_2]$ as desired. \square

4.3 An application to Herwig’s theorem

In its basic form, Herwig’s theorem [5, 7, 6] provides, for some given partial isomorphism p of a given finite relational structure \mathfrak{A} , an extension $\mathfrak{B} \supseteq \mathfrak{A}$ of \mathfrak{A} such that the given partial isomorphism p of \mathfrak{A} extends to a full automorphism of \mathfrak{B} . It generalises a corresponding theorem by Hrushovski [9], which makes the same assertion about graphs. Elegant combinatorial proofs of both theorems can be found in the paper by Herwig and Lascar [7], which also generalises them further to classes of structures that avoid certain weak substructures (cf. Corollary 4.16 below). The variant in which *every* partial isomorphism of the original finite structure extends to an automorphism of the extension is more

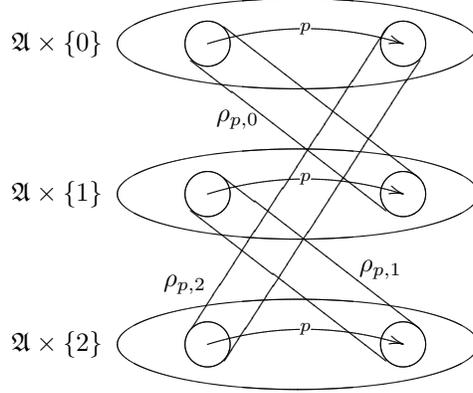


Figure 6: Basic idea for Herwig's theorem/ \mathbb{Z}_3

useful for some purposes, but is also an immediate consequence of the basic form, via simple finite iteration. W.l.o.g. we restrict attention to relational structures with a single relation R of some fixed arity r .

Theorem 4.10 (Herwig's Theorem). *For every finite relational structure $\mathfrak{A} = (A, R^{\mathfrak{A}})$ there is a finite extension $\mathfrak{B} = (B, R^{\mathfrak{B}}) \supseteq \mathfrak{A}$ such that every partial isomorphism of \mathfrak{A} is the restriction of some automorphism of \mathfrak{B} .*

Note that the term 'extension' as applied here means that \mathfrak{A} is an *induced substructure* of \mathfrak{B} , denoted $\mathfrak{A} \subseteq \mathfrak{B}$, which means that $A \subseteq B$ and $R^{\mathfrak{A}} = R^{\mathfrak{B}} \cap A^r$. A partial isomorphism of \mathfrak{A} is a partial map on A , $p: \text{dom}(p) \rightarrow \text{image}(p)$ that is an isomorphism between $\mathfrak{A} \upharpoonright \text{dom}(p)$ and $\mathfrak{A} \upharpoonright \text{image}(p)$ (induced substructures).

We here reproduce Herwig's theorem in an argument based on groupoidal constructions, which may offer a starting point for further generalisations. Before that, we prove the version of Herwig's theorem for a single partial isomorphism p from scratch, with a construction that will motivate the new proof of the full version.

Let $\mathfrak{A} = (A, R^{\mathfrak{A}})$ be a finite R -structure, p a partial isomorphism of \mathfrak{A} . We first provide a canonical infinite solution to the extension task for p and \mathfrak{A} . Let $\mathfrak{A} \times \mathbb{Z} = (A \times \mathbb{Z}, R^{\mathfrak{A} \times \mathbb{Z}})$ be the structure obtained as the disjoint union of isomorphic copies of \mathfrak{A} , indexed by \mathbb{Z} . Let \approx be the equivalence relation over $A \times \mathbb{Z}$ that identifies (a, n) with $(p(a), n + 1)$ for $a \in \text{dom}(p)$; we think of \approx as induced by partial matchings or local bijections

$$\begin{aligned} \rho_{p,n}: \text{dom}(p) \times \{n\} &\longrightarrow \text{image}(p) \times \{n + 1\} \\ (a, n) &\longmapsto (p(a), n + 1). \end{aligned}$$

Let \approx be the equivalence relation induced by the $\rho_{p,n}$, so that, for $m \leq n$,

$$(a_1, m) \approx (a_2, n) \text{ iff } a_2 = p^{n-m}(a_1).$$

The interpretation of R in $\mathfrak{A}_\infty := (\mathfrak{A} \times \mathbb{Z})/\approx$ is

$$R^{\mathfrak{A}_\infty} := \{[\bar{a}, m]: \bar{a} \in R^{\mathfrak{A}}, m \in \mathbb{Z}\},$$

where $[\bar{a}, m]$ is shorthand for the tuple of the equivalence classes of the components (a_i, m) w.r.t. \approx . By construction, \mathfrak{A} is isomorphic to the induced substructure represented by the slice $\mathfrak{A} \times \{0\} \subseteq \mathfrak{A} \times \mathbb{Z}$, on which \approx is trivial: $(a, 0) \approx (a', 0) \Leftrightarrow a = a'$. Since p and the $\rho_{p,n}$ are partial isomorphisms, the quotient w.r.t. \approx does not induce new tuples in the interpretation of R that are represented in the slice $\mathfrak{A} \times \{0\}$.

The shift $n \mapsto n - 1$ in the second component induces automorphisms $\eta: (a, n) \mapsto (a, n - 1)$ and $\eta: [a, n] \mapsto [a, n - 1]$ of $\mathfrak{A} \times \mathbb{Z}$ and of $(\mathfrak{A} \times \mathbb{Z})/\approx$. The automorphism η of $(\mathfrak{A} \times \mathbb{Z})/\approx$ extends the realisation of p in $\mathfrak{A} \times \{0\}$, since for $a \in \text{dom}(p) \subseteq A$:

$$\eta([a, 0]) = [a, -1] = [p(a), 0].$$

Therefore $\mathfrak{B}_\infty := (\mathfrak{A} \times \mathbb{Z})/\approx$ is an infinite solution to the extension task.

It is clear that the domain $\text{dom}(p^k)$ of the k -fold composition of the partial map p is eventually stable, and, for suitable choice of k , also induces the identity on $\text{dom}(p^k)$. Fixing such k , we look at the correspondingly defined quotient

$$\mathfrak{B} := (\mathfrak{A} \times \mathbb{Z}_k)/\approx,$$

which, by the same reasoning, embeds \mathfrak{A} isomorphically in the induced substructure represented by the slice $\mathfrak{A} \times \{0\}$, and has the automorphism $\eta: [a, n] \mapsto [a, n - 1]$ (now modulo k), which extends p . So \mathfrak{B} is a finite solution to the extension task.

We turn to the extension task for a specified collection P of partial isomorphisms of $\mathfrak{A} = (A, R^{\mathfrak{A}})$. Let $S := (\mathbb{Z}_3)^P$ be the set of P -indexed tuples over $\{0, 1, 2\}$. For $p \in P$ we denote as $s \mapsto s +_p 1$ the cyclic successor map of \mathbb{Z}_3 in application to the p -th component, which fixes all other components. Consider the vertex set $V := A \times S = A \times (\mathbb{Z}_3)^P$, partitioned according to the natural projections into subsets $V = \bigcup_s V_s$, where $V_s = A \times \{s\}$. We expand V to an R -structure $\mathfrak{H} = (V, R^{\mathfrak{H}})$ in the natural manner, so that we may think of \mathfrak{H} as a collection of disjoint copies of \mathfrak{A} . With $p \in P$ we associate partial bijections $\rho_{p,s}$ from V_s to $V_{s+_p 1}$

$$\begin{aligned} \rho_{p,s}: \text{dom}(p) \times \{s\} &\longrightarrow \text{image}(p) \times \{s+_p 1\} \\ (a, s) &\longmapsto (p(a), s+_p 1). \end{aligned}$$

Putting $E = \{(s, s+_p 1), (s+_p 1, s): s \in S, p \in P\}$, this structure also gives rise to an I -graph for $I = (S, E)$,

$$H = (V, (V_s), (R_e)),$$

where $R_{(s, s+_p 1)}$ is the graph of $\rho_{p,s}$, $R_{(s+_p 1, s)}$ its converse. Then the following are symmetries of H and its expansion to the R -structure \mathfrak{H} :

$$\begin{aligned} \eta_p: V &\longrightarrow V \\ (a, s) &\longmapsto (a, s -_p 1), \end{aligned}$$

whose induced symmetry of I maps $s \in S = (\mathbb{Z}_3)^k$ to $s+p1$.

Let $\hat{\mathfrak{A}}$ be a realisation of the overlap pattern specified by H that has all the symmetries as described in Corollary 4.9: in particular it lifts the automorphisms η_p of H to automorphisms of $\hat{\mathfrak{A}}$ and has vertical automorphisms that relate any two hyperedges that project to the same V_s .

Since any two partition subsets V_s of H are related by an automorphism of H , any two hyperedges of $\hat{\mathfrak{A}}$ are related by an automorphism of $\hat{\mathfrak{A}}$. We call such a hypergraph homogeneous:

Definition 4.11. We call a hypergraph (B, S) *homogeneous* if its automorphism group acts transitively on its set of hyperedges: for $s, s' \in S$ there is an automorphism η of (B, S) such that $\eta(s) = s'$. For a subgroup Aut_0 of the full automorphism group of (B, S) , we say that (B, S) is *Aut₀-homogeneous* if any two hyperedges in S are related by an automorphism from that subgroup.

We expand $\hat{\mathfrak{A}}$ to an R -structure $\mathfrak{B} := (\hat{A}, R^{\mathfrak{B}})$ by lifting $R^{\hat{\mathfrak{A}}}$ from every V_s to every hyperedge \hat{s} of $\hat{\mathfrak{A}}$ that projects to V_s through $\pi_{\hat{s}}: \hat{s} \rightarrow V_s$, putting

$$R^{\mathfrak{B}} = \bigcup_{\hat{s} \in \hat{S}} \{\bar{a} \in \hat{s}^r : \pi_{\hat{s}}(\bar{a}) \in R^{\hat{\mathfrak{A}}}\}.$$

We claim that

- (i) $\mathfrak{B} \upharpoonright \hat{s} \simeq \mathfrak{A}$ for every hyperedge \hat{s} of $\hat{\mathfrak{A}}$;
- (ii) the lift of the η_p to automorphisms of $\hat{\mathfrak{A}}$ are automorphisms also of the R -structure \mathfrak{B} ; so are the ‘vertical’ automorphisms of $\hat{\mathfrak{A}}$ that relate hyperedges that project to the same V_s ;
- (iii) up to suitable ‘vertical’ automorphisms $\zeta_{p, \hat{s}}$, η_p expands the realisation of p in $\pi_{\hat{s}}: \mathfrak{B} \upharpoonright \hat{s} \simeq \mathfrak{A}$:

$$\begin{array}{ccc} \hat{s} & \xrightarrow{\zeta_{p, \hat{s}} \circ \eta_p} & \hat{s} \\ \pi_{\hat{s}} \downarrow & & \downarrow \pi_{\hat{s}} \\ A \times \{s\} = V_s & \xrightarrow{p \times \{s\}} & V_s = A \times \{s\} \end{array}$$

For (i), it is clear that \mathfrak{A} is isomorphic to a weak substructure of $\mathfrak{B} \upharpoonright \hat{s}$; the point of the assertion is that $\mathfrak{B} \upharpoonright \hat{s}$ does not acquire new R -tuples through the identification between overlaps. But condition (i) in the definition of realisations, Definition 3.15, precisely concerns this restriction: the full overlap between any two distinct hyperedges is generated by a composition of partial matchings, which in our case are partial isomorphisms of the R -structure.

For (ii) we observe that the automorphisms of $\hat{\mathfrak{A}}$, which account for its homogeneity, are all compatible with the projections $\pi_{\hat{s}}: \hat{s} \rightarrow V_s = A \times \{s\}$ and even with their compositions with the projections to A :

$$\begin{array}{ccc} \pi_{\hat{s}}^A: \hat{s} & \longrightarrow & A \\ \hat{a} & \longmapsto & a \text{ if } \pi_{\hat{s}}(\hat{s}) = (a, s). \end{array}$$

This is clear for the ‘vertical’ automorphisms of $\hat{\mathfrak{A}}$, since they are compatible with $\pi_{\hat{s}}$ (cf. Corollary 4.9); it is also true for the automorphisms induced by the η_p , because the η_p themselves commute with the projections of the V_s to A .

For (iii), finally, it suffices to look at the operation of η_p in restriction to that part of \hat{s} that corresponds to the domain of p , i.e., to $(\pi_{\hat{s}})^{-1}(\text{dom}(p) \times \{s\}) = (\pi_{\hat{s}}^A)^{-1}(\text{dom}(p))$. For $\hat{a} \in (\pi_{\hat{s}}^A)^{-1}(\text{dom}(p))$, the image point $\hat{a}' := \eta^{\hat{\mathfrak{A}}}(\hat{a})$ satisfies

$$\hat{a}' \in \hat{s}' \quad \text{for some } \hat{s}' \text{ above } s' = s -_p 1 \quad \text{and s.t. } \pi_{\hat{s}'}^A(\hat{a}') = a.$$

Since $\hat{\mathfrak{A}}$ realises the overlap pattern specified by H , \hat{a}' is identified (via some $\rho_{p,s'}$ -edge) with an element of a hyperedge \hat{s}'' above s that projects to $p(a)$, i.e.,

$$\hat{a}' \in \hat{s}'' \quad \text{for some } \hat{s}'' \text{ above } s \quad \text{s.t. } \pi_{\hat{s}''}^A(\hat{a}') = p(a).$$

Therefore, composition of η_p with a ‘vertical’ automorphism of $\hat{\mathfrak{A}}$ that maps \hat{s}'' back to \hat{s} in a manner that is compatible with projections to V_s (and A), yields the desired extension of p .

This reproves Herwig’s theorem in a form that highlights the role of the hypergraph generated by automorphic images of \mathfrak{A} within the extension $\mathfrak{B} \supseteq \mathfrak{A}$.

Corollary 4.12 (Herwig’s Theorem). *For every finite $\mathfrak{A} = (A, R^{\mathfrak{A}})$ there is a finite relational structure $\mathfrak{B} = (B, R^{\mathfrak{B}})$ and an $\text{Aut}(\mathfrak{B})$ -homogeneous $|A|$ -uniform hypergraph (B, S) such that*

- (i) every automorphism of \mathfrak{B} is an automorphism of (B, S) ;
- (ii) $\mathfrak{A} \simeq \mathfrak{B} \upharpoonright s$ for all $s \in S$;
- (iii) every partial isomorphism of \mathfrak{B} whose domain and image sets are contained in hyperedges of (B, S) is induced by an automorphism of \mathfrak{B} .

We note that the statement as given can also be obtained as a corollary of Herwig’s theorem as stated in Theorem 4.10 above. Its new proof, however, allows for further variations w.r.t. the nature of the hypergraph (B, S) , which may for instance be required to be N -acyclic. Among other potential generalisations this reproduces the extension of Herwig’s theorem to the class of conformal structures and, e.g., of k -clique free graphs, obtained on the basis of Herwig’s theorem together with Lemma 3.4 in [8].

Here is one further consequence of the uniform construction in the new solution to the extension task. Consider \mathfrak{A} and P as above, together with the I -graph H on vertex set $V \times S$ for $S = (\mathbb{Z}_3)^P$ and $I = (S, E)$ as above, based on the $\rho_{s,p}$ for $p \in P, s \in S$. Let \mathfrak{B} be obtained, as above, by the natural pull-back of $R^{\mathfrak{A}}$ to a hypergraph $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ that realises the overlap pattern specified by the I -graph H and is sufficiently symmetric. In particular, for any $\hat{s} \in \hat{S}$, $\mathfrak{B} \upharpoonright \hat{s} \simeq \mathfrak{A}$, and for every $p \in P$ there is an $f_{\hat{s},p} \in \text{Aut}(\mathfrak{B})$ that extends $p_{\hat{s}}$, the copy of p in $\mathfrak{B} \upharpoonright \hat{s} \simeq \mathfrak{A}$.

Claim 4.13. *Let \mathfrak{B}' be any (finite or infinite) solution to the extension task for \mathfrak{A} and P , $\mathfrak{A} \subseteq \mathfrak{B}'$. Let $\hat{s}, \hat{s}' \in \hat{S}$ and fix an isomorphism $\pi: \mathfrak{B} \upharpoonright \hat{s} \simeq \mathfrak{A}$. Then*

the isomorphic embedding of $\mathfrak{B} \upharpoonright \hat{s}$ via π as $\mathfrak{A} \subseteq \mathfrak{B}'$ extends to a homomorphism $h: \mathfrak{B} \upharpoonright (\hat{s} \cup \hat{s}') \rightarrow \mathfrak{B}'$:

$$\begin{array}{ccc}
 & \mathfrak{B} \upharpoonright (\hat{s} \cup \hat{s}') & \\
 & \uparrow \subseteq & \searrow h \\
 \mathfrak{B} \upharpoonright \hat{s} & \xrightarrow{\pi} \mathfrak{A} & \xrightarrow{\subseteq} \mathfrak{B}'
 \end{array}$$

Proof. We observe that, according to the properties of a realisation, the overlap $\hat{s} \cap \hat{s}'$ in $\hat{\mathfrak{A}}$ is induced by a composition $\rho = \rho_w^H$ of partial maps ρ_{s_i, p_i} in H . The corresponding composition of partial isomorphisms p_i of \mathfrak{A} gives rise to a composition of automorphisms f'_{p_i} in \mathfrak{B}' , i.e., there is an automorphism $f'_w \in \text{Aut}(\mathfrak{B}')$ such that $f'_w(\mathfrak{A}) \cap \mathfrak{A}$ contains $\pi(\mathfrak{B} \upharpoonright (\hat{s} \cap \hat{s}')) \subseteq \mathfrak{A} \subseteq \mathfrak{B}'$. Therefore, a suitable isomorphism $\sigma: \mathfrak{B} \upharpoonright \hat{s}' \simeq \mathfrak{A} \subseteq \mathfrak{B}'$ can be composed with f'_w to produce $f'_w \circ \sigma: \mathfrak{B} \upharpoonright \hat{s}' \simeq f'_w(\mathfrak{A}) \subseteq \mathfrak{B}'$ compatible with the given isomorphism $\pi: \mathfrak{B} \upharpoonright \hat{s} \simeq \mathfrak{A} \subseteq \mathfrak{B}'$ on the overlap $\mathfrak{B} \upharpoonright (\hat{s} \cap \hat{s}')$. \square

If moreover the hypergraph template $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ in the construction of \mathfrak{B} is chosen to be N -acyclic, then this local extension process can be iterated to cover any substructure of \mathfrak{B} of size up to N , because any such substructure is acyclic and thus tree-decomposable with bags induced by hyperedges of $\hat{\mathfrak{A}}$. This gives the following.

Corollary 4.14. *For any finite R -structure \mathfrak{A} , any collection P of partial isomorphisms of \mathfrak{A} and for any $N \in \mathbb{N}$, there is a finite extension $\mathfrak{B} \supseteq \mathfrak{A}$ that satisfies the extension task for \mathfrak{A} and P and has the additional property that any substructure $\mathfrak{B}_0 \subseteq \mathfrak{B}$ of size up to N can be homomorphically mapped into any other (finite or infinite) solution \mathfrak{B}' to the extension task for \mathfrak{A} and P .*

Proof. Let \mathfrak{B} be obtained from a suitably symmetric and N -acyclic realisation $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ of the overlap pattern specified in H . Then any $\mathfrak{B}_0 \subseteq \mathfrak{B}$ of size up to N is acyclic and admits a tree decomposition $\lambda: T \rightarrow \{\hat{s} \cap B_0: \hat{s} \in \hat{S}\}$, where T is a finite directed tree s.t., for every $b \in B_0$, the set $\{v \in T: b \in \lambda(v)\}$ is connected in T . For an initial segment $T_i \subseteq T$, let $\lambda(T_i) := \bigcup\{\lambda(v): v \in T_i\} \subseteq B_0$.

We obtain the desired homomorphism $h: \mathfrak{B}_0 \rightarrow \mathfrak{B}'$ from a sequence of homomorphisms $h_i: \mathfrak{B}_0 \upharpoonright \lambda(T_i) \rightarrow \mathfrak{B}'$. Let the sequence of initial segments T_i of T start from the root and exhaust T by adding one new vertex v_i at a time, in the extension from T_i to T_{i+1} . The new node $v_i \in T_{i+1}$ covers

$$\mathfrak{B}_0 \upharpoonright \lambda(v_i) \subseteq \mathfrak{B}_0 \upharpoonright \hat{s}_i \subseteq \mathfrak{B} \upharpoonright \hat{s}_i \simeq \mathfrak{A}$$

for some $\hat{s}_i \in \hat{S}$; due to the connectivity condition, the intersection of $\lambda(v_i)$ with $\lambda(T_i)$ is fully contained in $\lambda(v)$ where v is the parent of v_i in T , so that $v \in T_i$. Hence an extension as in Claim 4.13 suffices to extend h_i to a suitable h_{i+1} . \square

From this analysis of the new construction for Herwig's theorem we therefore also obtain a major strengthening of Theorem 4.10 due to [7], which can be phrased as a finite-model property for the extension task. In [7], this finite model property is called EPPA for 'extension property for partial automorphisms'.

Definition 4.15. Let \mathcal{C} be a class of R -structures.

- (a) \mathcal{C} has the *finite model property* for the extension of partial isomorphisms to automorphisms (EPPA) if, for every finite $\mathfrak{A} \in \mathcal{C}$ and collection P of partial isomorphisms of \mathfrak{A} such that \mathfrak{A} has some (possibly infinite) extension $\mathfrak{A} \subseteq \mathfrak{B} \in \mathcal{C}$ in which each $p \in P$ extends to an automorphism, there is also a *finite* extension $\mathfrak{A} \subseteq \mathfrak{B}_0 \in \mathcal{C}$ in which each $p \in P$ extends to an automorphism.
- (b) \mathcal{C} is defined in terms of *finitely many forbidden homomorphisms* if, for some finite list of finite R -structures \mathfrak{C}_i , it consists of all R -structures \mathfrak{A} that admit no homomorphisms of the form $h: \mathfrak{C}_i \xrightarrow{\text{hom}} \mathfrak{A}$.

The following is now immediate from Corollary 4.14.

Corollary 4.16 (Herwig–Lascar Theorem). *Every class \mathcal{C} that is defined in terms of finitely many forbidden homomorphisms has the finite model property for the extension of partial isomorphisms to automorphisms (EPPA).*

4.4 Regular finite N -acyclic hypergraphs

Consider the example of the full 3-uniform hypergraph on a set of 4 vertices, $[4]^3$, which also arises as the boundary of the 3-simplex or the faces of the tetrahedron. This hypergraph is highly symmetric and the discussion in Sections 4.1 and 4.2 guarantees the existence of finite N -acyclic covers that preserve all these symmetries. In particular, the kind of covers obtained from reduced products with suitably symmetric I -groupoids for $I = I([4]^3) \simeq K_4$ (the intersection graph of $[4]^3$ is the 4-clique), produces finite N -acyclic covers whose automorphism group

- acts transitively on the set of hyperedges, and even on the set of adjacent pairs of hyperedges that share an edge,
- acts transitively by fibre-preserving automorphisms with every fibre of hyperedges, and
- realises every permutation within every hyperedge.

It shares these properties with the more geometric but just locally finite N -acyclic covers of $[4]^3$ that can be realised by simple tessellations of the hyperbolic plane.

Another potentially distinguishing feature of covers, at the purely combinatorial level, can be their branching. Locally finite as well as finite covers of $[4]^3$ that are conformal and 3-acyclic must be branched; any such cover will have overlaps of size 1 between hyperedges while such overlaps are not present in $[4]^3$. On the other hand, both kinds of covers of $[4]^3$ are faithful w.r.t. overlaps of size 2. This is obvious for the planar geometric arrangement. For the groupoidal construction it follows from the observation that only applications of the individual generators stemming from the original size 2 intersections produce intersection of size 2 in the cover; the sub-groupoids corresponding to the stabiliser of a given 2-subset of a hyperedge can only be copies of \mathbb{Z}_2 .

Other covers of $[4]^3$ can also have size 1 intersections between hyperedges that are not generated by a succession of size 2 overlaps that share and stabilise one single vertex. In fact, the glueing of two disjoint covers in two vertices that project to the same base vertex will always be a cover again, albeit one with an ‘unnecessary’ singularity.

It is not hard to see that $[4]^3$ admits a simple finite 5-acyclic cover, which locally at every vertex is a 2-fold cover; it can be obtained via identifications along the boundary of a hexagonal planar pattern of diameter 8. The intersection graph of this finite hypergraph does have edges corresponding to overlaps of size 1 – and these are generated by paths (of length 3 in this case) of successive transitions across intersections of size 2 that share just the vertex to be preserved. Towards further covers of this hypergraph, the generators corresponding to size 1 overlaps would therefore not be necessary in the groupoidal construction. This indicates that in general smaller sets of generators than the full edge set of the intersection graph may suffice for the groupoidal construction, and can be used to avoid some unnecessary singularities.

Here is just one immediate corollary of the above discussion.

Corollary 4.17. *For every $m, n, N \in \mathbb{N}$, the m -uniform hypergraph on n vertices, $[n]^m$ admits finite covers by m -uniform hypergraphs that are N -acyclic (conformal and without chordless cycles of lengths up to N). Such covers can be chosen to be faithful w.r.t. the overlaps corresponding to intersections of maximal size (viz. size $m - 1$), and symmetric, with an automorphism group that lifts every automorphism of $[n]^m$, acts transitively on the set of hyperedges, and by fibre-preserving automorphisms on each fibre of hyperedges.*

I expect that the new route to the construction of highly uniform and symmetric hypergraphs will find further applications and provide further insights.

Acknowledgement I thank Achim Blumensath for helpful comments on an earlier version of this technical report.

References

- [1] N. Alon. Tools from higher algebra. In R. Graham, M. Grötschel, and Lovász, editors, *Handbook of Combinatorics*, volume II, pages 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*, 30:497–513, 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*, volume 555, pages 56–62. 1989.
- [5] B. Herwig. Extending partial isomorphisms on finite structures. *Combinatorica*, 15:365–371, 1995.

- [6] B. Herwig. Extending partial isomorphisms for the small index property of many omega-categorical structures. *Israel Journal of Mathematics*, 107:93–124, 1998.
- [7] B. Herwig and D. Lascar. Extending partial isomorphisms and the profinite topology on free groups. *Transactions of the AMS*, 352:1985–2021, 2000.
- [8] I. Hodkinson and M. Otto. Finite conformal hypergraph covers and Gaifman cliques in finite structures. *Bulletin of Symbolic Logic*, 9:387–405, 2003.
- [9] E. Hrushovski. Extending partial isomorphisms of graphs. *Combinatorica*, 12:411–416, 1992.
- [10] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004.
- [11] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59 (1), 2012.