Finite Groupoids, Finite Coverings and Symmetries in Finite Structures

Martin Otto
Department of Mathematics
Technische Universität Darmstadt

March 2016

Abstract

We propose a novel construction of finite hypergraphs and relational structures that is based on reduced products with Cayley graphs of groupoids. To this end we construct groupoids whose Cayley graphs have large girth not just in the usual sense, but with respect to a discounted distance measure that contracts arbitrarily long sequences of edges within the same sub-groupoid (coset) and only counts transitions between cosets. Reduced products with such groupoids are sufficiently generic to be applicable to various constructions that are specified in terms of local glueing operations and require global finite closure. The main construction is applied to hypergraph coverings as well as to extension tasks that lift local symmetries to global automorphisms in the style of Herwig–Lascar extension properties for partial automorphisms.

mathematical subject classification: primary 20L05, 05C65, 05E18;
secondary 20B25, 20F05, 57M12

Summary. At the centre of this investigation are generic algebraic-combinatorial constructions of finite amalgams of finite structures. These amalgams are based on glueing instructions specified as free amalgamations between pairs of given structures. The input data for this problem correspondingly consist of a finite family of finite structures together with a set of partial isomorphism between pairs of these. They serve as local specifications of the overlap pattern between designated parts of a desired global solution or *realisation*. Such a realisation consists of a single finite structure made up of a union of designated substructures, each of which is isomorphically related to a member of the given family of templates, and with overlaps between these substructures in accordance with the given family of partial isomorphisms. The combinatorial core of the matter arises from the most fundamental instance of such a problem, where the input data consists of a family of just disjoint sets together with designated partial bijections between pairs of these sets. A realisation then is a hypergraph: its hyperedges represent the coordinate domains in an atlas of charts into the given family of sets whose changes of coordinates stem from the given family of partial bijections between them. We provide generic constructions of such finite realisations that allow us to respect all intrinsic symmetries of the specification and to achieve any degree of local acyclicity for the global intersection pattern between the constituent substructures. These constructions are based on reduced products with suitable finite groupoids. The existence of such finite groupoids, which need to have strong acyclicity properties w.r.t. cosets, constitutes one of the new results of this paper. More specifically, a construction that interleaves amalgamation arguments for Cayley graphs with groupoidal actions establishes the following.

Theorem 0.0. *For every $N \in \mathbb{N}$, every finite set of generators can be expanded into a finite groupoid that is N-acyclic in the strong sense that it admits no cyclic pattern of length up to N of non-trivial intersections between cosets formed w.r.t. sub-groupoids generated by subsets of the given set of generators.*

For specific further constraints in terms of induced inverse semigroup actions and in terms of symmetries see Theorem 2.21 and Corollary 4.5. The main combinatorial result uses reduced products with such groupoids and concerns the existence of finite realisations (cf. Theorem 3.24 and Corollary 4.10).

Theorem 0.1 (**Main Theorem**). *Every abstract finite specification of an overlap pattern between disjoint sets, by means of partial bijections between them, admits a realisation in a finite hypergraph. Moreover, for any $N \in \mathbb{N}$, this realisation can be chosen to be N-acyclic, i.e., such that every induced sub-hypergraph of up to N vertices is acyclic, and to preserve all symmetries of the given specification.*

As one a key application of the main theorem we look at finite realisations of the overlap pattern of a given finite hypergraph, and thus obtain finite branched coverings of any desired local degree of hypergraph acyclicity in the following sense (cf. Theorem 3.23).

Theorem 0.2. *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.*

As a key application of model-theoretic importance we obtain a new approach to extension properties for partial isomorphisms, which provide liftings of local symmetries to global symmetries (i.e., from partial to full automorphisms) in finite extensions, EPPA results in the terminology of Herwig and Lascar [11]. Our generic realisations of overlaps between copies of the given structure according to the given partial isomorphisms gives rise to such EPPA extensions. In fact they yield variant formulations of the powerful Herwig–Lascar EPPA theorem that are more specific w.r.t. to the local-to-global relationship between the parts and the whole and w.r.t. the symmetries involved (cf. Theorem 4.13 and Proposition 4.14).

Theorem 0.3. *For a finite relational structure $\mathfrak{A}$ and any collection P of partial isomorphisms of $\mathfrak{A}$ and $N \in \mathbb{N}$, there is a finite EPPA extension $\mathfrak{B} \supseteq \mathfrak{A}$ for $\mathfrak{A}$ and P such that any substructure $\mathfrak{B}_0 \subseteq \mathfrak{B}$ of size up to N maps homomorphically into any other (finite or infinite) EPPA solution for $\mathfrak{A}$ and P.*

Main contributions and connections. At the algebraic and combinatorial level, the constructions presented here pave the way to interesting new examples of finite groupoids and of associated Cayley graphs. In both cases, the main technical challenge and the important novelty in our constructions lies in the specific acyclicity properties, i.e., in a very strong control over cyclic configurations of bounded lengths. This control far exceeds the scope of classical Cayley graph constructions as discussed in [1] and may provide further classes of interesting examples in combinatorial and algorithmic contexts where local acyclicity of decompositions matters. The construction of reduced products with the Cayley graphs of finite groupoids that satisfy the appropriate finitary acyclicity properties is shown to be a natural tool to obtain suitable realisations and also to lend itself directly to the construction of finite coverings of controlled acyclicity for hypergraphs. The resulting hypergraph coverings provide interesting classes of highly homogeneous and highly acyclic finite hypergraphs, and thus synthetic examples of interest in relation to structural decomposition techniques in combinatorics and finite and algorithmic model theory. Methodologically, acyclicity of hypergraphs has long been recognised as an important criterion in combinatorial, algorithmic and logical contexts [5, 4]. In the well understood setting of graphs, local acyclicity can be achieved in finite bisimilar unfoldings or coverings based on Cayley groups of large girth, which is the classical notion of finitary acyclicity, [16]. We here see that Cayley groupoids (rather than groups) and the much stronger notion of coset acyclicity (rather than just large girth) together support the adequate generalisations in the hypergraph setting. Applications of closely related structural transformations to questions in logic, and especially in finite model theory, have already been explored, e.g., in [12, 17, 19, 8] w.r.t. expressive completeness results as well as w.r.t. algorithmic issues [2, 3] for families of guarded logics. The genericity of our local-to-global constructions based on suitable groupoids, as exemplified in their applications to EPPA extension issues which considerably exceed those in [12], also points to promising further applications in the study of automorphisms of countable structures built from finite substructures and of amalgamation classes that arise in the model-theoretic and algebraic analysis of homogeneous structures [15]. The tools and techniques proposed here may thus provide new and interesting examples of symmetry phenomena in finite as well as in countable structures. At the more conceptual level, the groupoidal constructions presented here point to interesting discrete and even finite analoga of classical concepts like branched coverings [7] and local symmetries [14] inviting further exploration.

Organisation of the paper. Section 1 amplifies on the ideas sketched above, by way of a broader exposition including simple examples that highlight some of the technical challenges. Section 2 introduces I-graphs and I-groupoids and their Cayley graphs and deals with the construction of finite I-groupoids with coset acyclicity properties (Theorem 0.0 above). Section 3 concerns the construction of realisations via reduced products of I-graphs with suitable I-groupoids. The main theorems on realisations and on coverings (Theorems 0.1 and 0.2 above) are proved in Section 3. Readers interested in coverings, and not in the more general setup used for realisations of abstract overlap specifications, can focus on Sections 3.3, 3.5 and 3.6 as far as the theme of finite coverings is concerned. Section 4 explores symmetries in realisations and applies them to the extension of local to global symmetries; this in particular yields EPPA extensions as in Theorem 0.3 above in a natural application of our generic constructions to overlap patterns between copies of the given structure.

1 Introduction

Consider a partial specification of some global structure by descriptions of its local constituents and of the links between these, in terms of allowed and required direct overlaps between pairs of local constituents. Such specifications typically have generic, highly regular, free, infinite realisations in the form of tree-like acyclic objects. We here address the issue of finite realisations, which should meet similar combinatorial criteria in terms of genericity and symmetry; instead of full acyclicity, which is typically unattainable in finite realisations, we look for specified degrees of acyclicity.

Overlaps can be specified by partial bijections between the local constituents. It turns out that the universal algebraic and combinatorial properties of groupoids, which can be abstracted from the composition behaviour of partial bijections, support a very natural approach to the construction of certain highly symmetric finite instances of hypergraphs and relational structures that provide the desired finite realisations. We use hypergraphs as abstractions for the decomposition of global structures into local constituents. As a collection of subsets of a given structure, the collection of hyperedges specifies a notion of locality: the local view comprises one hyperedge at a time. Depending on context, individual hyperedges may carry additional local structure, e.g., interpretations of relations stemming from an underlying relational structure. The analogy with atlases of charts for manifolds is a good guide: regarding the hyperedges as local coordinate neighbourhoods for charts, which provide local descriptions of the target object (the desired realisation), the overlap specification contains all the information about the available coordinatisations (in the individual charts) and about the inverse semigroup or pseudo-group of coordinate changes (for overlaps between charts).

Realisations. We deal with the synthesis task seeking a realisation as a global object that conforms to given local data, where the local data concern descriptions of the constituent pieces (to be implemented via charts that bijectively relate them to the given templates) and their overlap, which manifest themselves as changes of coordinates between the local views. We offer a versatile and generic solution to the finite synthesis problem posed by such local specifications of hypergraphs or of relational structures. Depending on the data even the very existence of a *finite* solution may not be immediate. In other circumstances, if some finite realisation is explicitly given or has been obtained in a first step, we aim to meet additional global criteria in special, qualified finite realisations. The global criteria under consideration are

(i) criteria concerning controlled acyclicity w.r.t. the natural gradation of hypergraph acyclicity or tree decomposability for finite structures;

(ii) criteria of global symmetry in the sense of a rich automorphism group that extends rather than breaks the symmetries of the given specification.

In the most general case, the specification of the overlap pattern to be realised comes as a disjoint family of abstract regions $(V_s)_{s \in S}$ together with a collection of partial bijections $(\rho_e)_{e \in E}$, where each $e \in E$ has specified source and target sites $s, s' \in S$ and ρ_e is a partial bijection between sites V_s and $V_{s'}$. We think of the set S as an index set for the different types of local constituents; and of E as an index set for the types of pairwise overlaps. Formally, E will be the set of edges in a multi-graph with vertex set S.

Figure 1 provides an example of a possible overlap specification between two sites V_s and $V_{s'}$ along with two different modes of overlap, e_1 and e_2. Note that minimal requirements w.r.t. an isomorphic (i.e., bijective) embedding of V_s into a desired realisation makes it necessary that, in this example, the e_1-overlap and the e_2-overlap of one and the same copy of V_s need to go to different copies of $V_{s'}$, and similarly every $V_{s'}$-copy will have to overlap with distinct V_s-copies.

Globally, a realisation of an overlap specification consists of partly overlapping components $V_{\hat{s}}$ that are each isomorphic to one of the local patches V_s via some local projection $\pi_{\hat{s}} \colon V_{\hat{s}} \to V_s$, and such that every copy of V_s overlaps with copies of suitable $V_{s'}$ according to the overlap specified in the partial bijections ρ_e for those edges $e \in E$ that link s to some s'. Figure 2 gives a local impression of how a partial bijection ρ_e between sites V_s and $V_{s'}$ of the overlap specification is to be realised as an actual overlap of isomorphic copies of these patches V_s

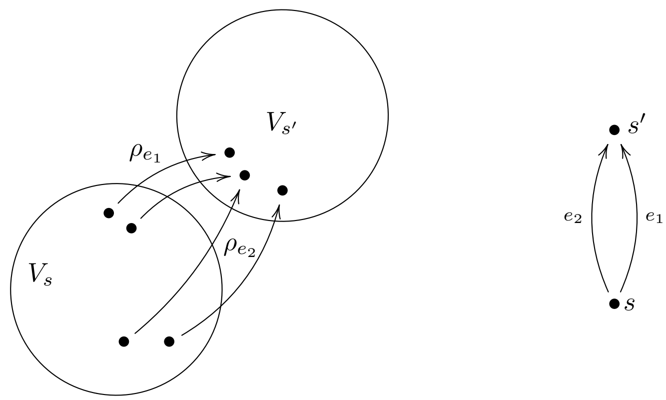

Figure 1: Links between two sites in overlap specification and in the underlying incidence pattern.

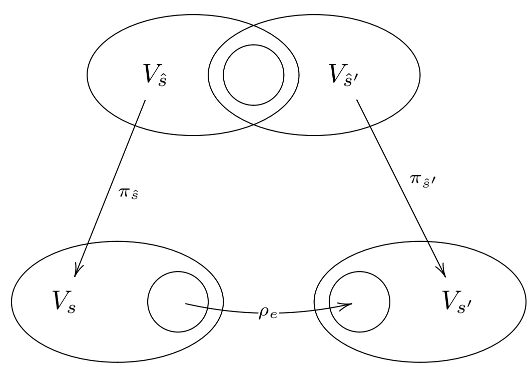

Figure 2: A realisation of some ρ_e as an overlap at one locus of a realisation.

and $V_{s'}$ within a realisation.

As discussed in [14], groupoids and inverse semigroups naturally occur in connection with the description of global structure by means of local coordinates. The systems of partial bijections that represent changes between different local coordinates can be abstracted as either inverse semigroups (sometimes also called pseudo-groups) or as groupoids (ordered groupoids in the terminology of [14]). As suggested above, the use of groupoids in the construction of realisations may be associated with atlases for a global realisation whose local views and overlaps between local views are specified in the given overlap pattern. The local sites $(V_s)_{s \in S}$ of the specification form the coordinate domains of an atlas for the realisation, in which local patches $V_{\hat{s}}$ corresponding to V_s are isomorphically related to their template V_s via local projections $\pi_{\hat{s}}$; the projections serve as the local charts, the ρ_e as changes of coordinates.

Coverings. More concrete overlap specifications arise from an actual structure composed of local patches. A non-trivial further realisation task in this context asks for a replication of the given overlap pattern in another finite structure, with extra constraints on the global properties of the new realisation. Intuitively, one can think of a finite process of partial unfolding. Formally, we cast this as a covering problem at the level of hypergraphs. These may be the hypergraphs induced by some notion of locality in other kinds of structures; the whole approach thus naturally extends to such settings. A hypergraph covering aims to reproduce the overlap pattern 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. The conceptual connections with topological notions of (branched) coverings [7] are apparent, but we keep in mind that here we insist on finiteness so that full acyclicity as in universal coverings cannot generally be expected.

Graph coverings. It may be instructive to compare first the situation for graphs rather than hypergraphs. Here the covering would be required to provide lifts for every edge – and, by extension, every path – in the base graph, at every vertex in the covering graph above a given vertex in the base graph. For an unbranched covering, the in- and out-degrees of a covering node would also be required to be the same as for the corresponding vertex in the base graph. And indeed, graphs do allow for unbranched finite coverings in this sense, which achieve any desired finite degree of acyclicity (i.e., can avoid cycles of length up to N for any desired threshold N), as the following result from [16] shows.

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

The situation for hypergraphs is more complicated because the separation between local and global aspects is blurred compared to graphs. This is due to the fact that a succession of transitions from one hyperedge to the next through

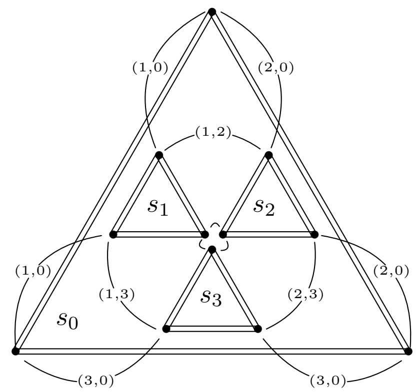

Figure 3: Overlap specification for the facets of the 3-simplex.

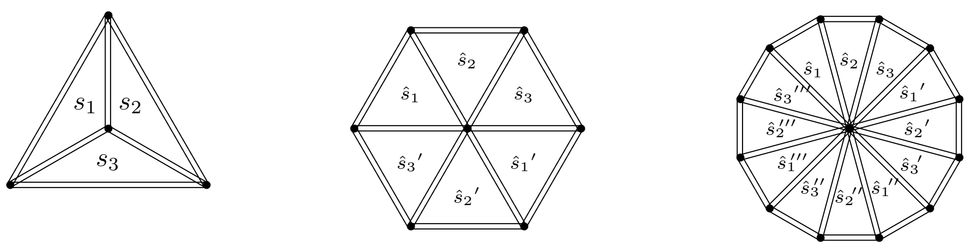

Figure 4: Local unfoldings around one vertex in the 3-simplex.

some partial overlaps typically preserves some vertices while exchanging others; in a graph on the other hand, two successive edge transitions cannot maintain the presence of any vertex. Correspondingly, even the notion of N-acyclicity for hypergraphs is somewhat non-trivial.

Hypergraph coverings. Consider a covering task for the simple hypergraph consisting of the facets of the 3-simplex. This hypergraph is associated with the faces of the tetrahedron or it may be seen as just the complete 3-uniform hypergraph on four vertices. A representation of the overlap specification between the hyperedges (faces) is provided in Figure 3. Possible local views around a single vertex of a finite covering hypergraph are indicated in Figure 4. It is clear that even locally, cycles cannot be avoided (in any trivial sense of avoiding cycles): the cyclic succession of overlapping copies of the hyperedges s_1, s_2 and s_3 around a single shared vertex will have to close back onto itself in any finite covering. This also means that the incidence degrees of vertices with hyperedges of a certain type cannot be preserved in any non-trivial finite covering, or that branched rather than unbranched coverings will have to be considered.

Proposition 1.1 above shows that neither of these perceived obstacles, viz.,

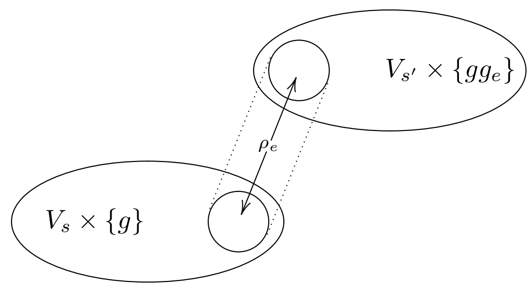

Figure 5: Identification between layers in a reduced product.

the lack of a clear local-vs-global distinction or the necessity to consider branched coverings, arises in the special case of graphs. The uniform and canonical construction of N-acyclic graph coverings according to Proposition 1.1, as given in [16], is based on a natural product between the given graph and 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 [6] in Alon's survey [1]. Some of these ideas were successfully lifted and applied to the construction of hypergraph coverings in [17]. For the combinatorial groups, the generalisation in [17] 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 finite 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 cosets formed by the stabilisers of one pivot and the next). But nothing as simple as a 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 [17] uses such Cayley groups only as one ingredient to achieve suitable hypergraph coverings. Further steps in the construction from [17] are no longer canonical; in particular, they do not preserve symmetries of the given hypergraph.

We here expand the amalgamation techniques that were explored for the combinatorial construction of highly acyclic Cayley graphs [17] 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 coverings as well as for the construction of finite hypergraphs according to other specifications. The new notion of Cayley groupoids allows for the construction of finite realisations of overlap specifications by means of natural reduced products with these groupoids. The basic idea of the use of groupoids in *reduced products* is the following. Think of a groupoid G whose elements are tagged by sort labels $s \in S$, which stand for the different sites V_s in the overlap specification; the groupoid is generated by elements $(g_e)_{e \in E}$ which link sort s to sort s' if the partial bijection ρ_e of the overlap specification links V_s to $V_{s'}$. Then a realisation of the overlap specification is obtained from a direct product that consists of disjoint copies $V_s \times \{g\}$ of V_s for every groupoid element g of sort s through a natural identification of elements in $V_s \times \{g\}$ and in $V_{s'} \times \{gg_e\}$ according to ρ_e, as sketched in Figure 5. This approach is highly canonical and supports realisations, and in particular also coverings, of far greater genericity and symmetry than previously available.

We address two main applications of this general construction. The first of these concerns the construction of coverings as a special case of a realisation task in which the overlap specification is induced by a given hypergraph (or a notion of locality in some decomposition of a given structure). See Theorem 0.2 above. The second main application concerns the extension of local symmetries to global symmetries in finite structures. This topic has been explored in model theoretic context in the work of Hrushovski, Herwig and Lascar [13, 9, 11]. As an application of our generic construction it also reflects on the role that pseudo-groups, inverse semigroups and groupoids play in the algebraic and combinatorial analysis of local symmetries, cf. [14]. In its basic form, due to Hrushovski [13], the statement is for graphs and says that every finite graph admits a finite extension such that every partial isomorphism within the given graph is induced by an automorphism of the extension. In other words, every local symmetry can be extended to a global symmetry in a suitable extension. Herwig [9] and then Herwig and Lascar [11] extended this combinatorial result first from graphs to arbitrary relational structures (hypergraphs), then to a conditional statement concerning finite extensions within a class $\mathcal{C}$ of relational structures provided there is an infinite extension of the required kind within $\mathcal{C}$. Herwig and Lascar use the term *extension property for partial automorphisms* (EPPA) for this extension task from local to global symmetries in finite structures. The classes $\mathcal{C}$ under consideration are defined in terms of forbidden homomorphisms. The most general constructions provided in [11] remain rather intricate, while the original construction for graphs from [13] was greatly simplified in [11]. We here derive a strong form of the Herwig–Lascar EPPA theorem as a natural application of the main theorem. See Theorem 0.3 above and Corollary 4.19. This application uses its full power w.r.t. its natural compatibility with symmetries (for a rich automorphism group) and control of cycles (to guarantee a consistent embedding of the given structure and the omission of forbidden homomorphisms).

2 Highly acyclic finite 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. In that construction, a subgroup of the full permutation group on the vertex set 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, 6]. 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. In other words, the R_e are partial matchings or the graphs of partial bijections within V. Then $e \in E$ induces a permutation g_e of the vertex set V, where g_e swaps the two vertices in every e-coloured edge. The $(g_e)_{e \in E}$ generate a subgroup of the group of all permutations of V. For suitable H, the Cayley graph induced by this group with generators $(g_e)_{e \in E}$ 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 'large girth' to a higher level of acyclicity. The second aspect is similar to the strengthening obtained in [17] 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 [17], we shall here construct groupoids and associated groupoidal Cayley graphs, which are edge- and vertex-coloured 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 some major technical challenges. It leads to objects that are better suited to hypergraph constructions than Cayley groups.

2.1 I-graphs

The basic idea for the specification of an overlap pattern was outlined in the introduction. We now formalise the concept with the notion of an *I-graph*. The underlying structure $I = (S, E)$, on which the notion of an I-graph will depend, is a multi-graph structure whose vertices $s \in S$ label the available sites and whose edges $e \in E$ label overlaps between these sites. This structure I serves as the *incidence pattern* for the actual overlap specification in I-graphs that instantiate sites and overlaps by concrete sets V_s for $s \in S$ and partial bijections ρ_e for $e \in E$.

Definition 2.1. An *incidence pattern* is a finite directed multi-graph $I = (S, E)$ with edge set $E = \dot{\bigcup}_{s,s' \in S} E[s, s']$, where $e \in E[s, s']$ is an edge from s to s' in I, with an involutive, fixpoint-free[1] edge reversal $e \mapsto e^{-1}$ on E that bijectively maps $E[s, s']$ to $E[s', s]$.

[1] We do *not* identify e with e^{-1} even for loops $e \in E[s, s]$ of I.

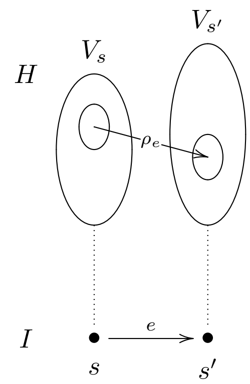

Figure 6: Local view of an I-graph H.

Remark 2.2. *There is an alternative, multi-sorted view, which may seem more natural from a categorical point of view. According to this view, an incidence pattern would be a structure* $I = (S, E, \iota_0, \iota_1, (\ )^{-1})$ *with two sorts* S *and* E*, where* $\iota_0, \iota_1 \colon E \to S$ *specify the start- and endpoints of edges* $e \in E$*, so that* $e \in E[\iota_0(e), \iota_1(e)]$*. Here edge reversal corresponds to simply swapping* ι_0 *and* ι_1*.*

The multi-sorted view will be our guide when we shall discuss symmetries of incidence patterns in Section 4. For notational convenience we stick to the shorthand format $I = (S, E)$ but keep in mind that this notation suppresses the two-sorted picture, the typing of the edges, and the operation of edge reversal.

Definition 2.3. 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 strictly partial[2] bijection ρ_e from V_s to $V_{s'}$, and such that the R_e (the ρ_e) are compatible with edge reversal, i.e., $R_{e^{-1}} = (R_e)^{-1}$ (or $\rho_{e^{-1}} = \rho_e^{-1}$).

In the following we use interchangeably the functional terminology of partial bijections ρ_e and the relational terminology of partial matchings R_e: it will be convenient to pass freely between the view of an I-graph H as either a coloured graph or as a family of disjoint sets linked by partial bijections.

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 graphs, where reflexive e-edges (loops) are allowed if e is a loop in I.

[2] 'Strictly partial' is to say that $\mathrm{dom}(\rho_e) \subsetneq V_s$; this restriction is natural as it avoids redundancies in the form of inclusion relationships between patches. It also turns out to be essentially without loss of generality, as one could add an otherwise inactive new extra element s to each V_s that plays the role of a tag to make the sets V_s pairwise incomparable up to intended identifications.

2.1.1 Operation of the free I-structure

We discuss the structure of the set of all compositions of the partial bijections ρ_e in an I-graph. These partial bijections form an inverse semigroup [14], but our main emphasis is on other aspects. Firstly, the analysis of the composition structure of the ρ_e underpins the idea of I-graphs as specifications of overlap patterns and thus of our crucial concept of *realisations*. Secondly, it prepares the ground for the association with groupoids in Section 2.2.

For $I = (S, E)$, we let E^* stand for the set of all labellings of directed paths (walks) in I. A typical element of E^* 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 $1 \leqslant i \leqslant n$. We admit the empty labellings of paths of length 0 at $s \in S$, and distinguish them by their location s as λ_s.[3] The set E^* is partitioned into subsets $E^*[s,t]$, which, for $s, t \in S$, consist of the labellings of paths from s to t in I, so that in particular $\lambda_s \in E^*[s,s]$. For $w = e_1 \dots e_n \in E^*[s,t]$, we write $w^{-1} := e_n^{-1} \dots e_1^{-1}$ for the *converse* in $E^*[t,s]$, which is obtained by reverse reading w and replacement of each edge label e by its reversal e^{-1}. The set E^* carries a partially defined associative concatenation operation

$$(w, w') \in E^*[s,t] \times E^*[t,u] \quad \longmapsto \quad ww' \in E^*[s,u],$$

which has the empty word $\lambda_s \in E^*[s,s]$ as the neutral element of sort s. One may think of this structure as a groupoidal analogue of the familiar word monoids, but we do not regard e^{-1} as an inverse to e (there is no cancellation of factors ee^{-1} in E^*). For further reference, we denote it as the *free I-structure*

$$\mathfrak{I}^* = (E^*, (E^*[s,t])_{s,t \in S}, \cdot, (\lambda_s)_{s \in S}).$$

Consider an I-graph $H = (V, (V_s), (R_e))$. The partial bijections ρ_e prescribed by the relations R_e together with their compositions along paths in E^* induce a structure of the same type as $\mathfrak{I}^*$, in fact a natural homomorphic image of $\mathfrak{I}^*$ as follows. For $e \in E[s,s']$, let ρ_e be the partial bijection between V_s and $V_{s'}$ induced by $R_e \subseteq V_s \times V_{s'}$. For $w \in E^*[s,t]$, define ρ_w as the partial bijection from V_s to V_t that is 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 $w \in E^*[s,t]$,

$$\rho_w \colon V_s \longrightarrow V_t$$

is a partial bijection, possibly empty. We obtain a homomorphic image of the free I-structure $\mathfrak{I}^* = (E^*, (E^*[s,t]), \cdot, (\lambda_s))$ under the map

$$\begin{array}{rcl} \rho \colon \mathfrak{I}^* & \longrightarrow & \big\{ f \colon f \text{ a partial bijection of } V \big\} \\ w = e_1 \dots e_n & \longmapsto & \rho_w = \prod_{i=1}^n \rho_{e_i}. \end{array}$$

[3] 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; secondly, we distinguish empty words $\lambda_s \in E^*$, one for every $s \in S$.

Concatenation of paths/words maps to (partial) composition of partial maps:

$$\rho_{ww'} = \rho_{w'} \circ \rho_w$$

wherever defined, i.e., for $w \in E^*[s,t], w' \in E^*[t,u]$ so that $ww' \in E^*[s,u]$.

The converse operation $w \mapsto w^{-1}$ maps to the inversion of partial maps

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

Note that the domain of $(\rho_e)^{-1} \circ \rho_e$ is $\mathrm{dom}(\rho_e)$ and may be a proper subset of V_s. So we still do not have groupoidal inverses – this will be different only when we consider *complete* I-graphs (cf. Definition 2.5) in Section 2.2 below.

2.1.2 Coherent I-graphs

We isolate an important special class of *coherent* I-graphs in terms of a particularly simple, viz. 'untwisted', composition structure of the ρ_e. This concept involves a notion of global path-independence.

Definition 2.4. An I-graph $H = (V, (V_s), (R_e))$ is *coherent* if every composition ρ_w for $w \in E^*[s,s]$ is a restriction of the identity on V_s:

$$\rho_w \subseteq \mathrm{id}_{V_s} \quad \text{for all } s \in S, w \in E^*[s,s].$$

Note that coherence is a property of path-independence for the tracking of vertices via ρ_w along paths in I: $\rho_{w_1}(v) = \rho_{w_2}(v)$ for every pair $w_1, w_2 \in E^*[s,t]$ and for all $v \in V_s$ in $\mathrm{dom}(\rho_{w_1}) \cap \mathrm{dom}(\rho_{w_2})$. To see this, consider the path $w = w_1^{-1} w_2 \in E^*[t,t]$ and apply the map ρ_w, which is responsible for transport along this loop, to the vertex $\rho_{w_1}(v)$. Coherence may also be seen as a notion of *flatness* in the sense that the operation of the free I-structure does not twist the local patches in a non-trivial manner. The I-graph representations of hypergraphs to be discussed in Section 3.1 provide natural examples of coherent I-graphs. Of course I itself is trivially coherent when considered as an I-graph.

Coherence of H implies that overlaps between arbitrary pairs of patches V_s and V_t, as induced by overlaps along connecting paths in I, are well-defined, independent of the connecting path. We let $\rho_{st}(V_s) \subseteq V_t$ stand for the subset of V_t consisting of those $v \in V_t$ that are in the image of ρ_w for some $w \in E^*[s,t]$. Then $\rho_{st}(V_s) \subseteq V_t$ is bijectively related to $\rho_{ts}(V_t) \subseteq V_s$ by the partial bijection

$$\rho_{st} := \bigcup \bigl\{ \rho_w \colon w \in E^*[s,t] \bigr\},$$

which is well-defined due to coherence.

2.1.3 Completion of I-graphs

Definition 2.5. An I-graph H is *complete* if the R_e induce full rather than partial bijections, i.e., if, for all $e \in E[s,s']$, $\mathrm{dom}(\rho_e) = V_s$ and $\mathrm{image}(\rho_e) = V_{s'}$.

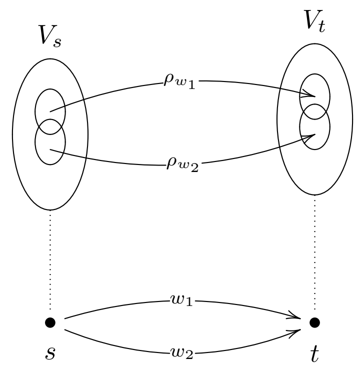

Figure 7: Coherence of I-graphs.

Complete I-graphs trivialise those complicating features of the composition structure of the ρ_e that arise from the partial nature of these bijections. Note that I itself may be regarded as a trivially complete I-graph. Cayley graphs of I-groupoids will be further typical examples of complete I-graphs, cf. Definition 2.11 below. A process of *completion* is required to prepare arbitrary given I-graphs for the desired groupoidal operation.

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:

$$\begin{array}{rcl} H \times I & := & \big(V \times S, (\tilde{V}_s), (\tilde{R}_e)\big) \text{ where, for } s \in S, \\ \tilde{V}_s & = & V \times \{s\}. \end{array}$$

For $e \in E[s, s']$, $s \neq s'$, the possibly incomplete R_e in H is lifted to $H \times I$ according to

$$\begin{array}{rcl} \tilde{R}_e & = & \big\{((v,s),(v',s')) \colon (v,v') \text{ an } e\text{-edge in } H\ \big\} \cup \\ & & \big\{((v',s),(v,s')) \colon (v,v') \text{ an } e\text{-edge in } H\ \big\} \cup \\ & & \big\{((v,s),(v,s')) \colon v \text{ not incident with an } e\text{-edge in } H\ \big\}; \end{array}$$

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

$$\begin{array}{rcl} \tilde{R}_e & = & \big\{((v,s),(v',s)) \colon (v,v') \text{ an } e\text{-edge in } H\ \big\} \cup \\ & & \big\{((v',s),(v,s)) \colon v/v' \text{ first/last vertex on a maximal } e\text{-path}^4 \text{ in } H\ \big\}. \end{array}$$

We note that this stipulation does indeed produce a complete I-graph: for $e \in E[s, s']$, it is clear from the definition of the $\tilde{R}_e$ that $\tilde{R}_e \subseteq \tilde{V}_s \times \tilde{V}_{s'}$ and that every vertex in $\tilde{V}_s$ has an outgoing e-edge and every vertex in $\tilde{V}_{s'}$ an incoming

[4] An e-path is a directed path in R_e^H.

e-edge; $\tilde{R}_e$ also is a bijection as required: in the non-reflexive case, 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 $\tilde{R}_{e^{-1}} = (\tilde{R}_e)^{-1}$ as required.

Observation 2.6. *If $H = (V, (V_s), (R_e))$ is a not necessarily complete I-graph, then $H \times I$ is a complete I-graph; the map*

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

embeds H isomorphically as a weak substructure. If I is loop-free, or if H is already complete w.r.t. to all loops $e \in E[s, s]$ of I for all $s \in S$, then H embeds into $H \times I$ as an induced substructure.

Proof. Note that the natural projection onto the first factor provides the inverse to σ on its image. Then $(v, v') \in R_e$ 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 $H \times I$. If $s \neq s'$ then $v \neq v'$ (as the V_s partition V) and (v, v') must be an e-edge of H. If $e \in E[s, s]$ is a loop of I, an e-edge $((v, s), (v', s))$ for $v, v' \in V_s$ may occur in $H \times I$ even though (v, v') is not an e-edge of H, but then v and v' were missing outgoing, respectively incoming, e-edges in H. $\square$

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

Definition 2.7. 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 $\sigma(V) = \{(v, s) \colon v \in V_s\}$. Identifying V with $\sigma(V) \subseteq H \times I$, we regard H as a weak subgraph of $\bar{H}$.

Corollary 2.8. *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, or complete cycles in H.

For the last claim observe that, for complete H, the vertex set of the isomorphic embedding $\sigma\colon H \to H \times I$ is closed under the edge relations $\tilde{R}_e$ 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 $H \times I$. $\square$

If $\alpha = \alpha^{-1} \subseteq E$ we write I_α for the reduct of I to its α-edges. We may 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 as well as I-graphs. Note that every I_α-graph is also an I-graph but for $\alpha \subsetneq E$ cannot be a complete I-graph. The α-reduct of the I-graph H is denoted $H \restriction \alpha$. Closures of subsets of I-graphs under α-edges (edges of colours $e \in \alpha$) will arise in some constructions below.

Lemma 2.9. *Let $\alpha = \alpha^{-1} \subseteq E$ and consider an I-graph H and its α-reduct $K := H \restriction \alpha$ as well as their completions $\bar{H}$ and $\bar{K}$ as I-graphs and the α-reducts $\bar{H} \restriction \alpha$ and $\bar{K} \restriction \alpha$ of those. Then $\bar{H} \restriction \alpha$ is an induced subgraph of $\bar{K} \restriction \alpha$,*

$$\bar{H} \restriction \alpha \subseteq \bar{K} \restriction \alpha,$$

and the vertex set of $\bar{H}$ is closed under α-edges within $\bar{K}$.

Proof. Recall that the completion $\bar{H} \subseteq H \times I$ consists of the connected component of the diagonal embedding $\sigma(V) \subseteq V \times S$ into $H \times I$. This connected component is formed w.r.t. the union of the edge relations $(\tilde{R}_e)_{e \in E}$ of $H \times I$. Similarly, the completion of K is formed by $\bar{K} \subseteq K \times I$, where the connected component of $\sigma(V) \subseteq V \times S$ is w.r.t. the union of the edge relations $(\tilde{R}'_e)_{e \in E}$ of $K \times I = (H \restriction \alpha) \times I$, for all $e \in E$. Let

$$D := \sigma(V) = \{(v, s) \colon s \in S, v \in V_s\},$$

which is the same set of seeds for the completions $\bar{H}$ and $\bar{K}$ as closures of D under $(\tilde{R}_e)_{e \in E}$ and $(\tilde{R}'_e)_{e \in E}$, respectively. For $e \in \alpha$, the edge relations $\tilde{R}_e$ and $\tilde{R}'_e$ of $H \times I$ and of $K \times I$ also coincide.

For $e \in E \setminus \alpha$, however, $\tilde{R}_e$ and $\tilde{R}'_e$ need not agree. Whenever (v, v') is an e-edge in H, for some $e \in E[s, s'] \setminus \alpha$, then this e-edge is not present in K, whence, for $s \neq s'$,

$$\begin{array}{lll} ((v,s),(v',s')),((v',s),(v,s')) \in \tilde{R}_e & \text{in } H \times I & \text{(double arrows in Fig. 8)} \\ ((v,s),(v,s')),((v',s),(v',s')) \in \tilde{R}'_e & \text{in } K \times I & \text{(single arrows in Fig. 8).} \end{array}$$

For $e \in E[s, s] \setminus \alpha$, no relevant discrepancies occur, since additional e-edges for $e \in E[s, s]$ have no effect on connectivity in either $H \times I$ or $K \times I$.

Since the vertices on the diagonal $(v, s), (v', s') \in D = \sigma(V)$ are vertices of both $\bar{H}$ and $\bar{K}$, the union of connected components that gives rise to $\bar{H}$ is included in the one that gives rise to $\bar{K}$: all four of the vertices (v, s), (v', s'), (v, s') and (v', s) are present in $\bar{K}$, while the off-diagonal pair (v, s'), (v', s) may or may not be present in $\bar{H}$. $\square$

2.2 I-groupoids

We obtain a groupoid operation on every complete I-graph H generated by the local bijections $\rho_e \colon V_s \to V_{s'}$ for $e \in E[s, s']$ as induced by the R_e. This step supports a groupoidal analogue of the passage from coloured graphs to Cayley groups.

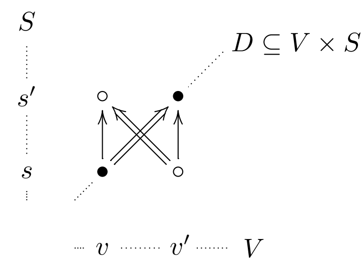

Figure 8: e-edges in $H \times I$ (double arrows) and $(H \restriction \alpha) \times I$ (single arrows) for $e \in E[s, s'] \setminus \alpha$, provided that (v, v') is an e-edge of H.

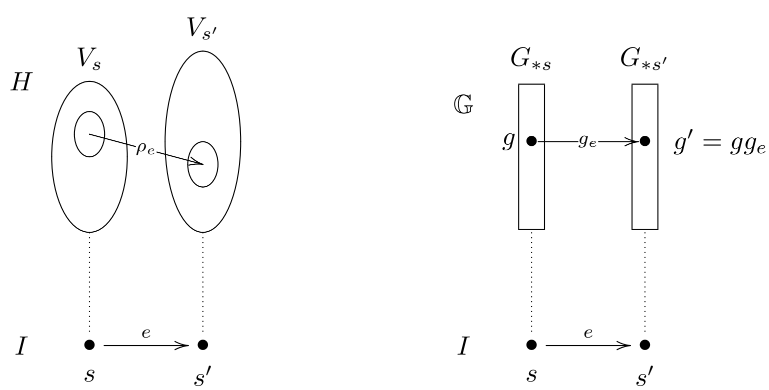

Figure 9: Local view of an I-graph H, and of an I-groupoid $\mathbb{G}$ as a complete I-graph: g_e bijectively links $G_{*s} = \bigcup_t G_{ts}$ and $G_{*s'} = \bigcup_t G_{ts'}$.

Definition 2.10. An S-*groupoid* is a structure $\mathbb{G} = \big(G, (G_{st})_{s,t \in S}, \cdot, (1_s)_{s \in S}\big)$ 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 $g^{-1} \cdot g = 1_t$.

We are looking to construct S-groupoids as homomorphic images of the free I-structure $\mathfrak{I}^*$ as discussed in Section 2.1.1. For the local view of an I-groupoid compare the right-hand side of Figure 9.

Definition 2.11. 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 $e_1 \dots e_n \in E^*[s, t]$.

An S-groupoid $\mathbb{G}$ that is generated by some family $(g_e)_{e \in E}$ for $I = (S, E)$ is called an I*-groupoid*.[5]

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

$$\begin{array}{rcl} {}^{\mathbb{G}} \colon \mathfrak{I}^* & \longrightarrow & \mathbb{G} \\ w = e_1 \dots e_n \in E^*[s,t] & \longmapsto & w^{\mathbb{G}} = \prod_{i=1}^{n} g_{e_i} \in G_{st}. \end{array}$$

The quotient of $\mathfrak{I}^*$ w.r.t. the equivalence relation induced by cancellation of converses, which identifies ee^{-1} with λ_s and $e^{-1}e$ with $\lambda_{s'}$ for $e \in E[s, s']$, provides a trivial example of an I-groupoid. 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' of I (these are not I-groupoids).

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

$$\mathbb{G}_\alpha := \mathbb{G} \restriction \{ w^{\mathbb{G}} \colon w \in \textstyle\bigcup_{st} \alpha^*_{st} \} \text{ with generators } (g_e)_{e \in \alpha}.$$

According to the above, $\mathbb{G}_\alpha$ may break up into separate and disjoint I_β-groupoids for the disjoint connected components I_β of I_α.

Recall from Section 2.1.1 how the free I-structure $\mathfrak{I}^*$ induces an operation on an I-graph H if we associate the partial bijections ρ_e of H with the generators

[5] It 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 $(e)_{e \in E}$.

$e \in E$ of $\mathfrak{I}^*$. The fixpoint-free edge reversal in I induces a converse operation $w \mapsto w^{-1}$ on $\mathfrak{I}^*$, which corresponds to inversion of partial bijections, $\rho_w \mapsto (\rho_w)^{-1} = \rho_{w^{-1}}$. In general this converse operation does not induce groupoidal inverses w.r.t. to the neutral elements $1_s = \mathrm{id}_{V_s}$: for $e \in E[s,s']$, the domain of $(\rho_e)^{-1} \circ \rho_e$ is $\mathrm{dom}(\rho_e)$, which may be a proper subset of V_s. It is the crucial distinguishing feature of complete I-graphs, cf. Definition 2.5, that we obtain the desired groupoidal inverse. If H is a complete I-graph, then $\rho_{w^{-1}} \circ \rho_w = (\rho_w)^{-1} \circ \rho_w = \mathrm{id}_{V_s}$ for any $w \in E^*[s,t]$, and the image structure obtained in this manner is an S-groupoid $\mathbb{G} =: \mathbb{G}(H)$:

$$\rho\colon \mathfrak{I}^* \quad\longrightarrow\quad \mathbb{G} =: \mathbb{G}(H),$$

where

$$\begin{aligned} \mathbb{G}(H) = \mathbb{G} &= \big(G, (G_{st})_{s,t\in S}, \cdot, (1_s)_{s\in S}\big), \\ G_{st} &= \{\rho_w \colon w \in E^*[s,t]\}. \end{aligned}$$

The groupoid operation $\cdot$ is the one imposed by the natural composition of members of corresponding sorts:

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

For $s \in S$, the identity $1_s := \mathrm{id}_{V_s}$ is the neutral element of sort G_{ss}, induced as $1_s = \rho_\lambda$ by the empty word $\lambda_s \in E^*[s,s]$. The natural groupoidal inverse is

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

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

Definition 2.12. For a complete I-graph H we let $\mathbb{G}(H)$ be the groupoid abstracted from H according to the above stipulations. We consider $\mathbb{G}(H)$ as an I-groupoid generated by $(\rho_e)_{e\in E}$. For a not necessarily complete I-graph H, let the I-groupoid $\mathbb{G}(H)$ be the I-groupoid $\mathbb{G}(\bar{H})$ induced by the completion $\bar{H}$ of H (cf. Definition 2.7).

It is easy to check that $\mathbb{G}(H)$ is an I-groupoid with generators $(\rho_e)_{e\in E}$ according to Definition 2.11. We turn to the analogue, for I-groupoids, of the notion of the Cayley graph.

Definition 2.13. 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 $V = G$,

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

and

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

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 $e \in E[s, s']$. Compare Figure 9.

Lemma 2.14. *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 $e \in E[s, s']$ this is the bijection

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

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 $\rho_w \colon V_s \to V_s$, 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 $\rho_w = \mathrm{id}_{V_s}$ as desired. $\square$

Remark 2.15. *The E-graphs of [17] and their role as Cayley graphs of groups are a special case of I-graphs, also in their roles as Cayley graphs of groupoids.*

In fact, an E-graph in the sense of [17] is a special I-graph for an incidence pattern of the form $I = (S, E)$ where S is a singleton set and E a collection of loops. An E-graph then is an I-graph in which every R_e is a partial matching. It follows that its completion consists of the symmetrisation of R_e augmented by reflexive edges at every vertex outside the domain and range of these matchings. For the induced I-groupoids abstracted from (complete) I-graphs consisting of matchings, this not only means that they are groups rather than groupoids, but also that they are generated by involutions, as in this case, $g_e = g_{e^{-1}}$.

2.3 Amalgamation of I-graphs

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

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 $\mathbb{G}_\alpha$-*coset* at g:

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

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 $\alpha' \subseteq \alpha$ 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 $\mathbb{G}_\alpha$.

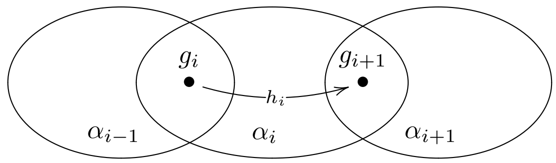

Figure 10: Amalgamation: overlap between cosets.

Suppose the I_α-graph H_α and the I_β-graph H_β are isomorphic to the Cayley graphs of sub-groupoids $\mathbb{G}_\alpha$ and $\mathbb{G}_\beta$, respectively. If $v_1 \in H_\alpha$ and $v_2 \in H_\beta$ are vertices of the same colour $s \in S$, 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 amalgam 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. This is a free amalgam over the identification of the two relevant $(\alpha\cap\beta)$-components. It is convenient to speak of (the Cayley graphs of) the sub-groupoids $\mathbb{G}_\alpha$ 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 that 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 $\mathbb{G}_{\alpha_i}$ 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 $(\mathbb{G}_{\alpha_i}, g_i, h_i, s_i)_{1\leqslant i\leqslant N}$ be a sequence of sub-groupoids with distinguished elements and vertex colours as indicated, and such that for all relevant i

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

For the last condition, compare Figure 10: it stipulates that g_i cannot be linked to g_{i+1} by an α_i-shortcut that merges the neighbouring α_{i-1}- and α_{i+1}-cosets within the α_i-coset that links g_i to g_{i+1}; intuitively, such a shortcut would allow us to eliminate entirely the step involving the α_i-coset.

If the above conditions are satisfied, then the pairwise amalgams

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

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

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

We call the free amalgam produced in this fashion *a chain of sub-groupoids* $\mathbb{G}_{\alpha_i}$ of length N. Condition (†) is important to ensure that the resulting structure is again an I-graph. Otherwise, an element of the critical intersection $g_i\mathbb{G}_{\alpha_{i-1}\alpha_i} \cap g_i h_i \mathbb{G}_{\alpha_i\alpha_{i+1}}$ could inherit new e-edges from both $\mathbb{G}_{\alpha_{i-1}}$ and from $\mathbb{G}_{\alpha_{i+1}}$, for $e \in (\alpha_{i-1} \cap \alpha_{i+1}) \setminus \alpha_i$.

2.4 Eliminating short coset cycles

Definition 2.16. A *coset cycle of length* n in an I-groupoid with generator set E is a sequence $(g_i)_{i\in\mathbb{Z}_n}$ of groupoid elements g_i (cyclically indexed) together with a sequence of generator sets (sets of edge colours) $\alpha_i = \alpha_i^{-1} \subseteq E$ 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 2.17. An I-groupoid is N*-acyclic* if it does not have coset cycles of length up to N.

We now aim for the construction of N-acyclic I-groupoids to be achieved in Theorem 2.21. The following definition of compatibility captures the idea that some I-groupoid $\mathbb{G}$ is at least as discriminating as the I-groupoid $\mathbb{G}(H)$ induced by the I-graph H.

Definition 2.18. 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 $w \in E^*[s,s]$,

$$w^{\mathbb{G}} = 1_s \text{ in } \mathbb{G} \quad \Longrightarrow \quad \rho_w = 1_s \text{ in } \mathbb{G}(H).$$

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

$$\begin{tikzcd}
\mathfrak{I}^* \arrow[d, "()^{\mathbb{G}}"'] \arrow[dr, "\rho"] & \\
\mathbb{G} \arrow[r, "\text{hom}"] & \mathbb{G}(H)
\end{tikzcd}$$

Compatibility of $\mathbb{G}$ with H also means that $\mathbb{G} = \mathbb{G}(\mathbb{G}) = \mathbb{G}(\mathbb{G} \,\dot{\cup}\, H)$ – and in this role, compatibility of sub-groupoids $\mathbb{G}_\alpha$ 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.

By definition, $\mathbb{G}(H)$ is compatible with H and $\bar{H}$ and, by Lemma 2.14, with its own Cayley graph.

Remark 2.19. *If K and H are any I-graphs, then $\mathbb{G}(H \,\dot{\cup}\, K)$ is compatible with K, $\bar{K}$ and with the Cayley graph of $\mathbb{G}(K)$.*

The following holds the key to avoiding short coset cycles. Note that only generator sets of even sizes are mentioned since we require closure under edge reversal.

Lemma 2.20. *Let $\mathbb{G}$ be an I-groupoid with generators $e \in E$, let $k, N > 0$, and assume that, for every $\alpha = \alpha^{-1} \subseteq E$ with $|\alpha| < 2k$, the sub-groupoid $\mathbb{G}_\alpha$ 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 $|\alpha| \leqslant 2k$, the sub-groupoid $\mathbb{G}^*_\alpha$ is compatible with chains $\mathbb{G}^*_{\alpha\beta_i}$ up to length N.

It will be important later that compatibility of $\mathbb{G}^*_\alpha$ with chains as in (ii) makes sure that $\mathbb{G}^*_\alpha$ cannot have cycles of cosets generated by sets $\alpha \cap \beta_i$ of length up to N: every such cycle in the groupoid $\mathbb{G}^*_\alpha = \mathbb{G}(\mathbb{G}^*_\alpha)$ would have to be a cycle also in the groupoid induced by any such chain, including those chains obtained as linear unfoldings of the proposed cycle (cf. Theorem 2.21 below).

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

Specifically, we let H 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 $\alpha = \alpha^{-1}, \beta_i = \beta_i^{-1} \subseteq E$, $1 \leqslant i \leqslant m \leqslant N$, where $|\alpha| \leqslant 2k$.

By construction and Remark 2.19, $\mathbb{G}^* = \mathbb{G}(\mathbb{G} \cup H)$ 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 (†) (page 22) the merged chain is isomorphic to $\mathbb{G}^*_\alpha$, 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 H. (That $\mathbb{G}^*$ is compatible with $\mathbb{G}_{\alpha'}$ is clear since $\mathbb{G}_{\alpha'}$ is itself a component of H and hence of $\mathbb{G} \cup H$; compatibility of $\mathbb{G}^*$ with $\mathbb{G}$ is obvious for the same reason.)

Consider then a component of the form $\bigoplus_{i=1}^{m}(\mathbb{G}_{\alpha\beta_i}, g_i, h_i, s_i)$. The α'-reducts of α'-connected components of its completion arise as substructures of the completions of merged chains of components of the form $\mathbb{G}_{\alpha'\alpha\beta}$, according to Lemma 2.9. Since $|\alpha'| < 2k$, the assumptions of the lemma imply compatibility of $\mathbb{G}_{\alpha'}$ with any such component. It follows that $\mathbb{G}^* = \mathbb{G}(\mathbb{G} \cup H)$ is compatible with all $\mathbb{G}_{\alpha'}$ for $|\alpha'| < 2k$, and thus $\mathbb{G}^*_{\alpha'} \simeq G_{\alpha'}$ for $|\alpha'| < 2k$. $\square$

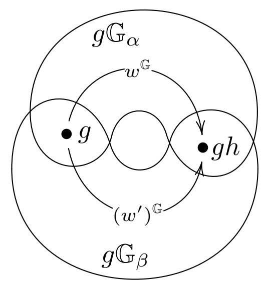

Figure 11: $w \in \alpha^*$ and $w' \in \beta^*$ form a coset 2-cycle unless g and gh are also linked in $(\alpha \cap \beta)^*$, i.e., unless $h \in \mathbb{G}_{\alpha\beta}$.

Theorem 2.21. *For every incidence pattern $I = (S, E)$ and $N \in \mathbb{N}$ there are finite N-acyclic I-groupoids with generators $e \in E$. Moreover, such an I-groupoid can be chosen to be compatible with any given I-graph H.*

Proof. Start from an arbitrary finite I-groupoid $\mathbb{G}_0$, or with $\mathbb{G}_0 := \mathbb{G}(H)$ in order to enforce compatibility with a given I-graph H. Then inductively apply Lemma 2.20 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 coset cycles of length up to N with cosets generated by sets of the form $\alpha\beta_i$ were $|\alpha| \leqslant 2k$. For $2k = |E|$, this rules out all coset cycles of length up to N. $\square$

Observation 2.22. *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 $w' \in \beta^*_{st}$ (cf. Figure 11). 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 coset 2-cycle with generator sets $\alpha_0 := \alpha$, $\alpha_1 := \beta$, unless the coset condition

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

of Definition 2.16 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 Reduced products with finite groupoids

3.1 Hypergraphs and hypergraph acyclicity

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 the set of vertices of $\mathfrak{A}$.

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 $E = \{(s, s') \colon s \neq s', s \cap s' \neq \emptyset\}$.

Note that the intersection graph $I(\mathfrak{A})$ captures the overlap pattern of the hyperedges of $\mathfrak{A}$. As an incidence pattern, $I(\mathfrak{A})$ suggests the representation of $\mathfrak{A}$ in an 'exploded view', which represents the hyperedges as disjoint sets with prescribed identifications between them for the overlaps. We shall come back to this representation of $\mathfrak{A}$ by an $I(\mathfrak{A}$-graph in Section 3.5.

The following criterion of hypergraph acyclicity is the natural and strongest notion of acyclicity (sometimes called α-acyclicity), cf., e.g., [5, 4]. It is in close correspondence with the algorithmically crucial notion of tree-decomposability (cf. Proposition 3.4 below) and with combinatorial notions of triangulation.

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

(i) conformality requires that every clique in the Gaifman graph $G(\mathfrak{A})$ is contained in some hyperedge $s \in S$;

(ii) chordality requires that every cycle in the Gaifman graph $G(\mathfrak{A})$ of length greater than 3 has a chord.

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

(iii) N-conformality requires that 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 requires that every cycle in the Gaifman graph $G(\mathfrak{A})$ of length greater than 3 and up to N has a chord.

N-acyclicity is a natural gradation or quantitative restriction of hypergraph acyclicity, in light of the following. Consider the induced sub-hypergraphs $\mathfrak{A} \restriction A_0$ of a hypergraph $\mathfrak{A} = (A, S)$, i.e., the hypergraphs on vertex sets $A_0 \subseteq A$ with hyperedge sets $S \restriction A_0 := \{s \cap A_0 \colon s \in S\}$. Then $\mathfrak{A}$ is N-acyclic if, and only if, every induced sub-hypergraph $\mathfrak{A} \restriction A_0$ for $A_0 \subseteq A$ of size up to N is acyclic.

Definition 3.3. A *tree decomposition* of a finite hypergraph (A, S) consists of an enumeration of the set S of hyperedges as $S = \{s_0, \dots, s_m\}$ such that for every $1 \leqslant \ell \leqslant m$ there is some $n(\ell) < \ell$ for which

$$s_\ell \cap \bigcup_{n<\ell} s_n \subseteq s_{n(\ell)}.$$

The important feature here is that $\bigcup S$ can be built up in a step-wise (and tree or forest-like) fashion, starting from the root patch s_0 by adding one member s_ℓ at a time in such a manner that the overlap of the new addition s_ℓ with the previous stage is fully controlled in the overlap with a single patch $n(\ell)$ in the previous stage. Thinking the whole process in reverse, we can see this as a decomposition process that reduces (A, S) to the empty hypergraph by simple retractions. For the following see [4].

Proposition 3.4. *A finite hypergraph (A, S) is acyclic if, and only if, it admits a tree decomposition. Correspondingly, it is N-acyclic if, and only if, every induced sub-hypergraph $\mathfrak{A} \restriction A_0$ for $A_0 \subseteq A$ of size $|A_0| \leqslant N$ admits a tree decomposition.*

3.2 Realisations of overlap patterns

The general case of an arbitrary I-graph over an arbitrary incidence pattern I seems to be a vast abstraction from the special case of an overlap pattern induced by the actual overlaps between hyperedges in an actual hypergraph. The notion of a realisation concerns this gap and formulates the natural conditions for a hypergraph to realise an abstract overlap specification. Intuitively, a realisation of an I-graph $H = (V, (V_s), (\rho_e))$ is a hypergraph to be viewed as an atlas: its hyperedges form the coordinate neighbourhoods for charts into the V_s with changes of coordinates in overlaps according to the ρ_e.

Definition 3.5. Let $I = (S, E)$ be an incidence pattern, $H = (V, (V_s), (R_e))$ an I-graph with induced partial bijections ρ_w between V_s and V_t for $w \in E^*[s, t]$. A hypergraph $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ is a *realisation* of the overlap pattern specified by H if there is a map $\pi \colon \hat{S} \to S$ and a matching family of bijections

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

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

$$\begin{tikzcd} & \hat{s} \cap \hat{s}' \arrow[dl, "\pi_{\hat{s}}"'] \arrow[dr, "\pi_{\hat{s}'}"] & \\ V_s \arrow[rr, "\rho_e"] & & V_{s'} \end{tikzcd} \qquad \begin{tikzcd} & \hat{s} \cap \hat{t} \arrow[dl, "\pi_{\hat{s}}"'] \arrow[dr, "\pi_{\hat{t}}"] & \\ V_s \arrow[rr, "\rho_w"] & & V_t \end{tikzcd}$$

Condition (i) says that all those local overlaps that should be realised according to H are indeed realised at corresponding sites in $\hat{\mathfrak{A}}$; condition (ii) says that all overlaps between hyperedges realised in $\hat{\mathfrak{A}}$ are induced by overlaps specified in H in a rather strict sense. In Section 3.5 below we shall look at a simpler concept of a *covering* of a given hypergraph $\mathfrak{A}$. Realisations of the overlap pattern $H(\mathfrak{A})$ abstracted from the given $\mathfrak{A}$, or its 'exploded view', will be seen to be special coverings. In this sense the notion of a realisation of an abstract overlap pattern (as specified by an I-graph) extends certain more basic notions of hypergraph coverings, in which the overlap pattern is specified by a concrete realisation.

Realisations and partial unfoldings. Regarding an I-graph H as a specification of an overlap pattern to be realised, it makes sense to modify H in manners that preserve the essence of that overlap specification. A natural idea of this kind would be to pass to a partial unfolding of H, which preserves the local links. We use the notion of a *bisimilar covering* at the level of the underlying incidence pattern for this purpose.

Let $\tilde{I} = (\tilde{S}, \tilde{E})$ and $I = (S, E)$ be incidence patterns. A homomorphism from $\tilde{I}$ to I is a map $\pi \colon \tilde{I} \to I$ respecting the (two-sorted) multi-graph structure so that, for $\tilde{e} \in \tilde{E}[\tilde{s}, \tilde{s}']$,

$$\pi(\tilde{e}) \in E[\pi(\tilde{s}), \pi(\tilde{s}')],$$

as well as the fixpoint-free involutive operations of edge reversal:

$$\pi(\tilde{e}^{-1}) = (\pi(\tilde{e}))^{-1}.$$

Definition 3.6. A surjective homomorphism $\pi \colon \tilde{I} \to I$ between incidence patterns $\tilde{I} = (\tilde{S}, \tilde{E})$ and $I = (S, E)$ is a *covering* of incidence patterns if it satisfies the following lifting property (known as the *back*-property in back&forth relationships like bisimulation equivalence):

$$(back)\colon \begin{cases} \text{for all } s \in S, e \in E[s, s'] \text{ and } \tilde{s} \in \pi^{-1}(s), \\ \text{there exists } \; \tilde{s}' \text{ and } \tilde{e} \in \tilde{E}[\tilde{s}, \tilde{s}'] \text{ s.t. } \pi(\tilde{e}) = e. \end{cases}$$

In the situation of the definition, an I-graph $H = \big(V, (V_s)_{s\in S}, (R_e)_{e\in E}\big)$ induces an $\tilde{I}$-graph $\tilde{H} = \big(\tilde{V}, (V_{\tilde{s}})_{\tilde{s}\in\tilde{S}}, (R_{\tilde{e}})_{\tilde{e}\in\tilde{E}}\big)$ on a subset $\tilde{V}$ of $V \times \tilde{S}$, where

$$\begin{aligned} &V_{\tilde{s}} := V_{\pi(\tilde{s})} \times \{\tilde{s}\} \;\subseteq\; \tilde{V} := \textstyle\bigcup_{\tilde{s}\in\tilde{S}} V_{\tilde{s}}, \\ &R_{\tilde{e}} := \big\{((v, \tilde{s}), (v', \tilde{s}')) \colon \tilde{e} \in \tilde{E}[\tilde{s}, \tilde{s}'], (v, v') \in R_e \text{ for } e = \pi(\tilde{e})\big\}. \end{aligned}$$

Lemma 3.7. *Suppose $\pi_0 \colon \tilde{I} \to I$ is a covering of the incidence pattern I, $\tilde{H}$ the $\tilde{I}$-graph induced by the I-graph H. Then every realisation of the overlap pattern specified by $\tilde{H}$ induces a realisation of the overlap pattern specified by H.*

Proof. Let $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ be a realisation of $\tilde{H} = (\tilde{V}, (V_{\tilde{s}})_{\tilde{s}\in\tilde{S}}, (R_{\tilde{e}})_{\tilde{e}\in\tilde{E}})$ with associated projections $\tilde{\pi} \colon \hat{S} \to \tilde{S}$ and $\tilde{\pi}_{\hat{s}} \colon \hat{s} \to V_{\tilde{\pi}(\hat{s})}$. Combining these projections with $\pi_0 \colon \tilde{S} \to S$ and the trivial projection π_1 that maps $V_{\tilde{s}} = V_{\pi_0(\tilde{s})} \times \{\tilde{s}\}$ to $V_{\pi_0(\tilde{s})}$, we obtain projections $\pi := \pi_0 \circ \tilde{\pi} \colon \hat{S} \to S$ and $\pi_{\hat{s}} := \pi_1 \circ \tilde{\pi}_{\hat{s}} \colon \hat{s} \to \pi_1(V_{\tilde{\pi}(\hat{s})}) = V_{\pi(\hat{s})}$, which allow us to regard $\hat{\mathfrak{A}}$ as a realisation of H. Towards the defining conditions on realisations, (i) is guaranteed by the *back*-property for π_0, while (ii) follows form the homomorphism condition for π_0 and the definition of $\tilde{H}$. In particular, $\mathrm{dom}(\rho_{\tilde{w}}) = \mathrm{dom}(\rho_w) \times \{\tilde{s}\}$ for any $\tilde{w} \in \tilde{E}^*[\tilde{s}, \tilde{t}]$ with projection $w = \pi_0(\tilde{w}) \in E^*[s, t]$ where $s = \pi_0(\tilde{s})$ and $t = \pi_0(\tilde{t})$. $\square$

3.3 Direct and reduced products with groupoids

Direct products. We define a natural *direct product* $H \times \mathbb{G}$ of an I-graph $H = (V, (V_s)_{s \in S}, (R_e)_{e \in E})$ with an I-groupoid $\mathbb{G}$. The construction may be viewed as a special case of a more general notion of a direct product between two I-graphs. In geometric-combinatorial terms we are interested in $H \times \mathbb{G}$ because it plays the role of a finite unfolding or covering of H, at least if $\mathbb{G}$ satisfies the compatibility condition in the sense of Definition 2.18.

For an I-graph H and I-groupoid $\mathbb{G}$ we define the direct product $H \times \mathbb{G}$ to be the following I-graph on the disjoint union of the products $V_s \times G_{*s}$, cf. Figure 12:

$$H \times \mathbb{G} := \Big(\bigcup_{s \in S} (V_s \times G_{*s}), \big(V_s \times G_{*s}\big)_{s \in S}, (R_e)_{e \in E} \Big)$$

where $R_e = \big\{ ((v,g), (\rho_e(v), gg_e)) \colon v \in V_s, g \in G_{*s} \big\}$ for $e \in E[s,s']$.

Just like H, this direct product admits an operation of the free I-structure in terms of compositions of partial bijections. These are based on the natural lifting of e to ρ_e (in H) and further to $\rho_e^{H \times \mathbb{G}}$ (in $H \times \mathbb{G}$) according to

$$\begin{array}{rcl} \rho_e^{H \times \mathbb{G}} \colon V_s \times G_{*s} & \longrightarrow & V_{s'} \times G_{*s'} \\ (v,g) & \longmapsto & (\rho_e(v), gg_e). \end{array}$$

This extends to paths $w \in E^*[s,t]$, which are lifted to ρ_w and further to $\rho_w^{H \times \mathbb{G}}$, cf. Figure 12.

Observation 3.8. *If $\mathbb{G}$ is compatible with H, then the liftings along different paths $w_1, w_2 \in E^*[s,t]$ linking the same groupoid elements agree in their common domains in $H \times \mathbb{G}$: for every $g \in G_{*s}$ and all $v \in \mathrm{dom}(\rho_{w_1}) \cap \mathrm{dom}(\rho_{w_2}) \subseteq V_s$,*

$$w_1^{\mathbb{G}} = w_2^{\mathbb{G}} \quad \Rightarrow \quad \rho_{w_1}(v,g) = \rho_{w_2}(v,g) \;\; \textit{in } H \times \mathbb{G}.$$

Proof. Wherever the composition $\rho_w^{H \times \mathbb{G}}$ is defined, it agrees in the first component with the operation of ρ_w on H and on the completion $\bar{H}$, which gives rise to $\mathbb{G}(H)$. Therefore $w_1^{\mathbb{G}} = w_2^{\mathbb{G}}$, or equivalently $(w_1^{-1} w_2)^{\mathbb{G}} = 1_s$, implies by compatibility that $\rho_{w_1^{-1} w_2} = \mathrm{id}_{V_s}$ in $\mathbb{G}(H)$, whence $\rho_{w_1}(v) = \rho_{w_2}(v)$ for $v \in \mathrm{dom}(\rho_{w_1}) \cap \mathrm{dom}(\rho_{w_2})$. $\square$

In other words, compatibility of $\mathbb{G}$ with H guarantees that any $w \in E^*[s,s]$ such that $w^{\mathbb{G}} = 1_s$ (i.e., any cycle in $\mathbb{G}$) induces a partial bijection

$$\rho_w \colon V_s \times G_{*s} \to V_s \times G_{*s}$$

that is compatible with the identity of $V_s \times G_{*s}$:

$$w^{\mathbb{G}} = 1_s \quad \Longrightarrow \quad \rho_w^{H \times \mathbb{G}} \subseteq \mathrm{id}_{V_s \times G_{*s}}.$$

The path-independence expressed in the observation is, however, characteristically weaker than coherence of $H \times \mathbb{G}$ as an I-graph, because we only compare paths that link the same groupoid elements. But $H \times \mathbb{G}$ also carries the structure of an $I(\mathbb{G})$-graph for the natural incidence pattern $I(\mathbb{G})$ associated with the Cayley graph of $\mathbb{G}$:

$$I(\mathbb{G}) = (G, \tilde{E}) \quad \text{where } \tilde{E} = \bigcup_{e \in E} \big\{ (g, gg_e) \colon e \in E[s,s'], g \in G_{*s} \big\}.$$

The relationship between $I(\mathbb{G})$ and I is that of a covering as in Definition 3.6 w.r.t. to the natural projection that maps $G_{*s} \subseteq G$ to $s \in S$. Moreover, $H \times \mathbb{G}$ is the $I(\mathbb{G})$-graph induced by H in the sense of Lemma 3.7. That lemma therefore tells us that any realisation of $H \times \mathbb{G}$ will provide a realisation of H. This in turn yields an interesting reduction of the general realisation problem to the realisation problem for coherent I-graphs, since the path-independence of Observation 3.8 precisely states that $H \times \mathbb{G}$ is coherent as an $I(\mathbb{G})$-graph in the sense of Definition 2.4. This route to realisations is pursued in Proposition 3.14.

Observation 3.9. *If $\mathbb{G}$ is compatible with H, then the direct product $H \times \mathbb{G}$ is coherent when viewed as an $I(\mathbb{G})$-graph. Every realisation of $H \times \mathbb{G}$ induces a realisation of H.*

Reduced products. The *reduced product* $H \otimes \mathbb{G}$ between an I-graph and an I-groupoid is simply obtained as the natural quotient of the direct product $H \times \mathbb{G}$ w.r.t. the equivalence relation $\approx$ induced by

$$(v, g) \approx (\rho_e(v), gg_e) \quad \text{for } e \in E[s,s'], v \in V_s, g \in G_{*s}.$$

Note that, by transitivity, for arbitrary $(v,g), (v',g') \in H \times \mathbb{G}$,

$$(v,g) \approx (v',g') \quad \text{iff } v' = \rho_w(v) \text{ for some } w \in E^*[s,t] \text{ with } g' = g \cdot w^{\mathbb{G}}.$$

We denote equivalence classes w.r.t. $\approx$ by square brackets, as in $[v,g] := \{(v',g') \colon (v',g') \approx (v,g)\}$ and extend this notation naturally to sets as in

$$[V_s, g] := [V_s \times \{g\}] = \{[v,g] \colon v \in V_s\}.$$

Definition 3.10. Let $H = (V, (V_s)_{s \in S}, (R_e)_{e \in E})$ be an I-graph, $\mathbb{G}$ an I-groupoid. The *reduced product* $H \otimes \mathbb{G}$ is defined to be the hypergraph $H \otimes \mathbb{G} := \hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ with vertex set

$$\hat{A} = \big\{ [v,g] \colon (v,g) \in \textstyle\bigcup_s V_s \times G_{*s} \big\}$$

and set of hyperedges

$$\hat{S} = \big\{ [V_s, g] \colon s \in S, g \in G_{*s} \big\},$$

where square brackets denote passage to equivalence classes w.r.t. $\approx$ as indicated above.

Note that the hyperedges $[V_s, g]$ of $H \otimes \mathbb{G}$ are induced by the patches V_s of H. Hyperedges $[V_s, g]$ and $[V_{s'}, g']$ for $(s,g) \neq (s',g')$ are always distinct as sets because all the partial bijections ρ_e are required to be strictly partial according to our definition of I-graphs (cf. Definition 2.3).

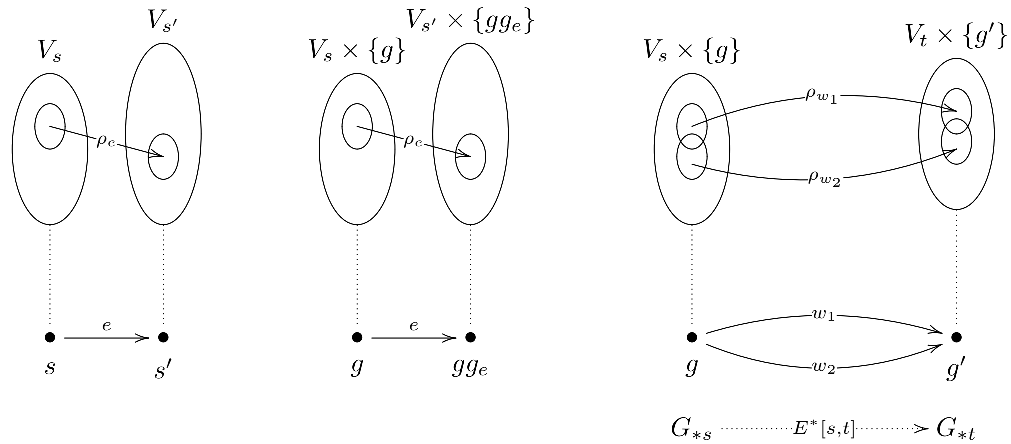

Figure 12: Fibres in an I-graph and in its product with an I-groupoid.

Lemma 3.11. *If $\mathbb{G}$ is compatible with the I-graph $H = (V, (V_s), (E_s))$, then the natural projection*

$$\begin{array}{rcl} \pi_{s,g} \colon [V_s, g] & \longrightarrow & V_s \\ [v, g] & \longmapsto & v \end{array}$$

is well-defined in restriction to each hyperedge $[V_s, g]$ of $H \otimes \mathbb{G}$, and relates each hyperedge $[V_s, g] = \{[v, g] \colon v \in V_s\}$ for $g \in G_{*s}$ bijectively to V_s.

Proof. It suffices to show that $[v, g] = [v', g]$ implies $v = v'$, which shows that $\pi_{s,g}$ is well-defined. By compatibility of $\mathbb{G}$ with H, $w^{\mathbb{G}} = 1$ implies $\rho_w \subseteq \mathrm{id}_{V_s}$, for any $w \in E^*[s, s]$. (For the last step compare Observation 3.8 about path-independence.) □

An even higher degree of path-independent transport in $H \times \mathbb{G}$ is achieved if H itself is a coherent I-graph in the sense of Definition 2.4 and if $\mathbb{G}$ is at least 2-acyclic in the sense of Definition 2.17.

Recall from the discussion after Definition 2.4 that coherence implies the existence of a unique and well-defined partial bijection ρ_{st} between those elements of V_s and V_t that can be linked by any ρ_w for $w \in E^*[s, t]$, viz.,

$$\rho_{st} = \bigcup \{\rho_w \colon w \in E^*[s, t]\}.$$

Observation 3.12. *If H is coherent and $\mathbb{G}$ is compatible with H and 2-acyclic, then there is, for any $g \in G_{*s}$ and $g' \in G_{*t}$ a unique maximal subset of V_s among all subsets*

$$\mathrm{dom}(\rho_w) \subseteq V_s \quad \textit{for those } w \in E^*[s, t] \textit{ with } g' = g \cdot w^{\mathbb{G}}.$$

Hence, in the reduced product $H \otimes \mathbb{G}$, the full intersection between hyperedges $[V_s, g]$ and $[V_t, g']$ is realised by the identification via ρ_w for a single path $w \in E^*[s,t]$ such that $g' = g \cdot w^{\mathbb{G}}$.

The second formulation is the key to the importance of this observation towards the construction of realisations.

The observation could also be phrased in terms of the direct product as follows. For any fixed site $V_s \times \{g\}$ in $H \times \mathbb{G}$, the maximal overlap of $V_s \times \{g\}$ with any other site $V_t \times \{g'\}$ via some composition of partial bijections $\rho_e^{H\times\mathbb{G}}$ is well-defined. In the reduced product $H \otimes \mathbb{G}$, this maximal overlap represents the full intersection between the hyperedges $[V_s, g]$ and $[V_t, g']$.

Proof of the observation. It suffices to show that, if $w_1, w_2 \in E^*[s,t]$ are such that $w_1^{\mathbb{G}} = w_2^{\mathbb{G}}$, then $\mathrm{dom}(\rho_{w_1}) \cup \mathrm{dom}(\rho_{w_2}) \subseteq \mathrm{dom}(\rho_w)$ for some suitable choice of $w \in E^*[s,t]$ for which also $w^{\mathbb{G}} = w_i^{\mathbb{G}}$.

By coherence of H, the image of any $v \in \mathrm{dom}(\rho_{w_i}) \subseteq V_s$ under any applicable composition of partial bijections ρ_e is globally well-defined, so that the point-wise image of the sets $\mathrm{dom}(\rho_{w_i}) \subseteq V_s$ at any V_u can be addressed as $\rho_{su}(\mathrm{dom}(\rho_{w_i})) \subseteq V_u$. As ρ_{w_i} maps every element of $\mathrm{dom}(\rho_{w_i})$ along the path w_i, this path can only involve generators from

$$\alpha_i = \bigcup_{u,u' \in S} \big\{ e \in E[u,u'] \colon \rho_{su}(\mathrm{dom}(\rho_{w_i})) \subseteq \mathrm{dom}(\rho_e) \big\},$$

for $i = 1, 2$. So $w_i^{\mathbb{G}} \in \mathbb{G}_{\alpha_i}$. By 2-acyclicity of $\mathbb{G}$, $g_1 := 1_s$ and $g_2 := w_2^{\mathbb{G}} = w_1^{\mathbb{G}}$ does not form a coset cycle w.r.t. the generator sets α_i. As $(g_1)^{-1} g_2 = w_2^{\mathbb{G}} = w_1^{\mathbb{G}} \in \mathbb{G}_{\alpha_1}$ and $(g_2)^{-1} g_1 = (w_2^{\mathbb{G}})^{-1} \in \mathbb{G}_{\alpha_2}$, it follows that the coset condition must be violated. So there must be some

$$w \in (\alpha_1 \cap \alpha_2)^* \text{ with } w^{\mathbb{G}} = w_1^{\mathbb{G}} = w_2^{\mathbb{G}}.$$

But

$$\alpha_1 \cap \alpha_2 = \bigcup_{u,u' \in S} \big\{ e \in E[u,u'] \colon \rho_{su}(\mathrm{dom}(\rho_{w_1}) \cup \mathrm{dom}(\rho_{w_2})) \subseteq \mathrm{dom}(\rho_e) \big\},$$

so that $w \in (\alpha_1 \cap \alpha_2)^*$ implies that $\mathrm{dom}(\rho_{w_i}) \subseteq \mathrm{dom}(\rho_w))$ for $i = 1, 2$ as desired. ☐

Corollary 3.13. *If $\mathbb{G}$ is compatible with the coherent I-graph $H = (V, (V_s), (E_s))$ and 2-acyclic, then the reduced product $H \otimes \mathbb{G}$ with its natural projections is a realisation of the overlap pattern specified in H.*

Proof. Compatibility of $\mathbb{G}$ with H guarantees that the natural projections

$$\begin{array}{rcl} \pi_{\hat{s}} \colon [V_s, g] & \longrightarrow & V_s \\ {[v, g]} & \longmapsto & v \end{array}$$

are well-defined in restriction to each hyperedge $[V_s, g]$ of $H \otimes \mathbb{G}$, and map this hyperedge bijectively onto V_s, by Lemma 3.11. For condition (i) in Definition 3.5, it is clear by construction of $H \otimes \mathbb{G}$ that for $e \in E^*[s, s']$ and $g \in G_{*s}$, $\hat{s} = [V_s, g]$ overlaps with $\hat{s}' := [V_{s'}, gg_e]$ according to ρ_e; this overlap cannot be strictly larger than $|\rho_e|$, as g_e is not in the sub-groupoid generated by $E \setminus \{e, e^{-1}\}$, due to 2-acyclicity of $\mathbb{G}$. Condition (ii) of Definition 3.5 is settled by Observation 3.12. $\square$

3.4 Realisations in reduced products

Combining the constructions of direct and reduced products with the existence of suitable groupoids we are ready to prove the first major step towards the main theorem on realisations: the existence of realisations for overlap patterns specified by arbitrary I-graphs. We shall then see in Section 3.6 below that the degree of acyclicity in realisations can be boosted to any desired level N through passage to N-acyclic coverings. That will then take us one step closer, in Theorem 3.24, to the full statement of the main theorem, Theorem 0.1. For the full content of the theorem as stated there, however, compatibility with symmetries will have to wait until Section 4, see Corollary 4.10.

Proposition 3.14. *For every incidence pattern I and I-graph H, there is a finite hypergraph $\hat{\mathfrak{A}}$ that realises the overlap pattern specified by H.*

Proof. From a given I-graph H we first obtain its product $H \times \mathbb{G}$ with an I-groupoid $\mathbb{G}$ that is compatible with H (see Theorem 2.21 for existence). Regarding $H \times \mathbb{G}$ as a coherent $\tilde{I}$-graph for $\tilde{I} = I(\mathbb{G})$ (cf. Observation 3.9), we obtain a realisation of that $\tilde{I}$-graph $H \times \mathbb{G}$ by a reduced product with a 2-acyclic $\tilde{I}$-groupoid that is compatible with $H \times \mathbb{G}$, according to Corollary 3.13. As $\tilde{I} = I(\mathbb{G})$ is a covering of I in the sense of Definition 3.6 and $H \times \mathbb{G}$ the $\tilde{I}$-graph induced by H, Lemma 3.7 guarantees that the resulting hypergraph is indeed a realisation of H. $\square$

We remark that the approach to realisations as presented above is different from the one outlined in [18, 19]. That construction first produces some kind of 'pre-realisations' $H \otimes \mathbb{G}$, which satisfy condition (i) for realisations but have more identifications than allowed by condition (ii). These pre-realisations can then be modified in a second unfolding step w.r.t. a derived incidence pattern to set condition (ii) right. The present approach seems more natural in that it puts realisations rather than coverings and unfoldings at the centre.

3.5 Coverings by reduced products

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

Definition 3.15. A hypergraph homomorphism $h \colon \hat{\mathfrak{A}} \to \mathfrak{A}$ between the hypergraphs $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ and $\mathfrak{A} = (A, S)$ is a *hypergraph covering* (of $\mathfrak{A}$ by $\hat{\mathfrak{A}}$) if it satisfies the *back*-property w.r.t. hyperedges: for every $\hat{s} \in \hat{S}$, $s = h(\hat{s}) \in S$ and $s' \in S$ there is some $\hat{s}' \in \hat{S}$ such that $h(\hat{s}') = s'$ and $h(\hat{s} \cap \hat{s}') = s \cap s'$.

Exploded view. Definition 3.1 abstracts the intersection graph $I(\mathfrak{A}) = (S, E)$ where $E = \{(s, s') \colon s \neq s', s \cap s' \neq \emptyset\}$ from any hypergraph $\mathfrak{A} = (A, S)$. Regarding $I = I(\mathfrak{A})$ as an incidence pattern, we associate an I-graph $H(\mathfrak{A})$ with $\mathfrak{A}$, which describes the structural information in $\mathfrak{A}$ in an 'exploded view' as follows. It consists of the disjoint union of the hyperedges together with the identifications ρ_e induced by the overlaps $s \cap s'$ for $e = (s, s') \in E$, as follows:

$$H(\mathfrak{A}) = \big(V, (V_s)_{s \in S}, (R_e)_{e \in E}\big)$$

where $E = \{(s, s') \colon s \neq s', s \cap s' \neq \emptyset\}$ is the edge relation of the intersection graph, V is the disjoint union of the hyperedges $s \in S$,

$$V = \bigcup_{s \in S} V_s \quad \text{ where } V_s = s \times \{s\},$$

naturally partitioned into the $(V_s)_{s \in S}$, and with the R_e (or ρ_e) that identify overlaps in intersections:

$$R_e = \big\{((v, s), (v, s')) \colon v \in s \cap s'\big\} \quad \text{ for } e = (s, s') \in E.$$

Clearly $H(\mathfrak{A})$ is coherent. Of course $I(\mathfrak{A})$ and $H(\mathfrak{A})$ have several further special properties. E.g., $I(\mathfrak{A})$ is loop-free and a graph rather than a multi-graph, i.e., each non-trivial $E[s, s'] = \{(s, s')\} \subseteq E$ is a singleton set. Besides coherence of $H(\mathfrak{A})$, there is a strong transitivity property involving both $I(\mathfrak{A})$ and $H(\mathfrak{A})$: if $e = (s, s') \in E[s, s'], e' = (s', s'') \in E[s', s'']$ are such that $\rho_{e'} \circ \rho_e \neq \emptyset$, then $\rho_{e'} \circ \rho_e = \rho_{e''}$ for $e'' = (s, s'') \in E[s, s'']$. Through its intersection graph and exploded view, a hypergraph embodies prototypical instances of incidence patterns, of I-graphs as overlap specifications, as well as of realisations, since it also comes as its own trivial realisation in the sense of Definition 3.5.

But any realisation $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ of $H(\mathfrak{A})$ then is a hypergraph covering for $\mathfrak{A}$, albeit one that avoids certain redundancies w.r.t. intersections. Cf. property (ii) in Definition 3.5 for realisations for the following rendering of this condition for coverings, which we want to call *strict* coverings.[6]

Definition 3.16. A hypergraph covering $h \colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is *strict* if every intersection $\hat{s} \cap \hat{t}$ between hyperedges of $\hat{\mathfrak{A}}$ is induced by a sequence of intersections in $\mathfrak{A}$ in the sense that $h(\hat{s} \cap \hat{t}) \times \{\hat{s}\} = \mathrm{dom}(\rho_w)$ in $H(\mathfrak{A})$, for some path $w \in E^*[s, t]$ from $s = h(\hat{s})$ to $t = h(\hat{t})$ in $I(\mathfrak{A})$.[7]

Observation 3.17. *Every realisation of $H(\mathfrak{A})$ is a strict hypergraph covering w.r.t. the natural projection induced by the projections of the realisation.*

[6]Natural though the general notion of a hypergraph covering may be, it does not rule out, e.g., partial overlaps between different covers $\hat{s}, \hat{s}'$ of the same s; the example indicated in Figure 4 shows that some such branching behaviour can be unavoidable.

[7]Note that the map ρ_w in $H(\mathfrak{A})$ represents a composition of identities in intersections in terms of $\mathfrak{A}$ itself.

Proof. Let $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ with projections $\pi_0 \colon \hat{S} \to S$ and $\pi_{\hat{s}} \colon \hat{s} \to V_{\pi_0(\hat{s})}$ be a realisation of the overlap pattern specified by $H = H(\mathfrak{A})$. Writing π_1 for the projection to the first component in $V = \bigcup_{s \in S}(s \times \{s\})$, we obtain $\pi \colon \hat{\mathfrak{A}} \to \mathfrak{A}$ as the compositions

$$\pi := \textstyle\bigcup_{\hat{s}} \pi_1 \circ \pi_{\hat{s}}.$$

π is well-defined since condition (ii) for realisations makes sure that overlaps in $\hat{\mathfrak{A}}$ are induced by compositions ρ_w in H, which, in $H = H(\mathfrak{A})$, are trivial in composition with π_1:

$$\pi_1 \circ \rho_w \subseteq \mathrm{id}_A \text{ for any } w \in E^*.$$

It is easy to check that $\pi \colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is a hypergraph homomorphism (as $\pi_{\hat{s}}$ bijectively maps $\hat{s} \in \hat{S}$ onto $V_{\hat{s}} = \pi_0(\hat{s}) \times \{\pi_0(\hat{s})\}$) and satisfies the *back*-property (by condition (i) on realisations). ☐

As $H(\mathfrak{A})$ is coherent, every reduced product $H(\mathfrak{A}) \otimes \mathbb{G}$ of $H(\mathfrak{A})$ with a 2-acyclic groupoid $\mathbb{G}$ that is compatible with $H(\mathfrak{A})$, according to Corollary 3.13 provides a realisation, and hence a strict covering. The resulting reduced product can also be cast more directly as a natural reduced product with $\mathfrak{A}$ itself, which offers a more intuitive view, and puts fewer constraints on $\mathbb{G}$. The following essentially unfolds and combines the definitions of $H(\mathfrak{A})$ (cf. discussion below Definition 3.15) and $H \otimes \mathbb{G}$ (cf. Definition 3.10).

Definition 3.18. Let $\mathfrak{A} = (A, S)$ be a hypergraph, $\mathbb{G}$ an I-groupoid for $I = I(\mathfrak{A})$. The *reduced product* $\mathfrak{A} \otimes \mathbb{G}$ is the hypergraph $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ whose vertex set $\hat{A}$ is the quotient of the disjoint union of G_{*s}-tagged copies of all $s \in S$,

$$\hat{A} := \left(\textstyle\bigcup_{s \in S, g \in G_{*s}} s \times \{g\} \right) / \approx$$

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

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

Denoting the equivalence class of (a, g) as $[a, g]$, and lifting this notation to sets in the usual manner, the set of hyperedges of $\mathfrak{A} \otimes \mathbb{G}$ is

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

The covering homomorphism π is the natural projection $\pi \colon [a, g] \longmapsto a$.

We note that (a, g) is identified with (a, g') in this quotient if, and only if, $g' = g \cdot w^{\mathbb{G}}$ for some path $w \in E^*[s, t]$ such that $(a, s) \in \mathrm{dom}(\rho_w)$ (and hence $\rho_w(a, s) = (a, t)$) in $H = H(\mathfrak{A})$. We may think of the generators $e = (s, s') \in E_a$ as preserving the vertex a in passage from $a \in s$ to $a \in s'$, so that the g-tagged copy of s and the gg_e-tagged copy of s' are glued in their overlap $s \cap s'$.

It is easy to see directly that $\pi \colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is indeed a hypergraph covering.

Proposition 3.19. *For any hypergraph $\mathfrak{A}$ and I-groupoid $\mathbb{G}$, where $I = I(\mathfrak{A})$, the reduced product $\mathfrak{A} \otimes \mathbb{G}$ with the natural projection*

$$\begin{aligned} \pi\colon \mathfrak{A} \otimes \mathbb{G} &\longrightarrow \mathfrak{A} \\ [a, g] &\longmapsto a \end{aligned}$$

is a hypergraph covering. If $\mathbb{G}$ is 2-acyclic, then this covering is strict.

One also checks that, indeed, $\mathfrak{A} \otimes \mathbb{G} \simeq H(\mathfrak{A}) \otimes \mathbb{G}$.

Another useful link between realisations and coverings is the following.

Lemma 3.20. *If $\pi\colon \hat{\mathfrak{A}}' \to \hat{\mathfrak{A}}$ is a strict hypergraph covering of a realisation $\hat{\mathfrak{A}}$ of the overlap pattern specified by the I-graph H, then so is $\hat{\mathfrak{A}}'$, w.r.t. the projections induced by the natural compositions of those of the realisation $\hat{\mathfrak{A}}$ with π.*

Proof. Let $I = (S, E)$, $H = (V, (V_s), (R_e))$, $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ covered by $\hat{\mathfrak{A}}' = (\hat{A}', \hat{S}')$ through π, and let $\pi_0\colon \hat{S} \to S$ and $\pi_{\hat{s}}\colon \hat{s} \to V_{\pi_0(\hat{s})}$ be the projections through which $\hat{\mathfrak{A}}$ realises H. Then

$$\pi_0' := \pi_0 \circ \pi\colon \hat{S}' \to S \quad \text{together with} \quad \pi_{\hat{s}'} := \pi_{\pi(\hat{s}')} \circ \pi\colon \hat{s}' \to V_{\pi_0'(\hat{s}')}$$

serve as the projections required in the realisation of H by $\hat{\mathfrak{A}}'$. □

3.6 Acyclicity in coverings by reduced products

We show that the absence of short coset cycles in $\mathbb{G}$, i.e., N-acyclicity in the sense of Definition 2.17, implies corresponding degrees of hypergraph acyclicity in the sense of Definition 3.2 in reduced products $\mathfrak{A} \otimes \mathbb{G}$.

Remark: The analysis of cycles and cliques in the following two sections could be carried out in the slightly more general setting of realisations of *coherent* I-graphs in reduced products, rather than the setting of coverings by reduced products. We choose the latter for the sake of greater transparency. The difference is just that one would have to work with coherent translations of overlap regions $\rho_{st}(V_s) \subseteq V_t$ in $H \otimes \mathbb{G}$ instead of the much more intuitive use of lifted pre-images of the actual intersections between hyperedges in $\mathfrak{A}$ – this precisely is the advantage of having one realisation of $H(\mathfrak{A})$ already, albeit the trivial one by $\mathfrak{A}$ itself.

Chordality in coverings by reduced products

Lemma 3.21. *Let $\mathfrak{A} = (A, S)$ be a hypergraph, $\mathbb{G}$ an N-acyclic I-groupoid for $I := I(\mathfrak{A})$, the intersection graph of $\mathfrak{A}$. Then $\mathfrak{A} \otimes \mathbb{G}$ is N-chordal.*

Proof. Suppose that $([a_i, g_i])_{i \in \mathbb{Z}_n}$ is a chordless cycle in the Gaifman graph of $\mathfrak{A} \otimes \mathbb{G}$. W.l.o.g. the representatives (a_i, g_i) are chosen such that, for suitable

$s_i \in S$, $[s_{i+1}, g_{i+1}]$ is a hyperedge linking $[a_i, g_i]$ and $[a_{i+1}, g_{i+1}]$ 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

$$\alpha_i := E_{a_i} = \big\{(s, s') \in E \colon a_i \in s \cap s'\big\}$$

such that $g_{i+1} = g_i \cdot w^{\mathbb{G}}$. In particular, $g_i^{-1} g_{i+1} \in \mathbb{G}_{\alpha_i}$. We claim that $(g_i)_{i \in \mathbb{Z}_n}$ is a coset cycle w.r.t. the generator sets $(\alpha_i)_{i \in \mathbb{Z}_n}$, in the sense of Definition 2.16. 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 coset 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 $[a_{i-1}, g_{i-1}]$ and $[a_{i+1}, g_{i+1}]$ are linked by a chord induced by the hyperedge $[k, t]$:

(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 $[a_{i-1}, g_{i-1}] \in [t, k]$.

(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 $[a_{i+1}, g_{i+1}] \in [t, k]$.

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

Conformality in coverings by reduced products

Lemma 3.22. *Let $\mathfrak{A} = (A, S)$ be a hypergraph, $\mathbb{G}$ an N-acyclic I-groupoid for $I := I(\mathfrak{A})$, the intersection graph of $\mathfrak{A}$. Then $\mathfrak{A} \otimes \mathbb{G}$ is N-conformal.*

Proof. Suppose that $X := \{[a_i, g_i] \colon i \in n\}$ 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 $[t_i, k_i]$ such that $X_i := \{[a_j, g_j] \colon j \neq i\} \subseteq [t_i, k_i]$. Let $h_i := k_i^{-1} k_{i+1}$ and $\alpha_i := \bigcap_{j \neq i, i+1} E_{a_j}$. Note that $k_i^{-1} k_{i+1} \in \mathbb{G}_{\alpha_i}$. We claim that $(k_i)_{i \in \mathbb{Z}_n}$ with generator sets $(\alpha_i)_{i \in \mathbb{Z}_n}$ forms a coset cycle in $\mathbb{G}$ in the sense of Definition 2.16. 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 $X \subseteq [t, k]$ would follow.

Since $k \in k_i \mathbb{G}_{\alpha_i \alpha_{i-1}}$, and $[a_j, g_j] \in [t_i, k_i]$, i.e., $[a_j, g_j] = [a_j, k_i]$, for all $j \neq i$, clearly $[a_j, g_j] \in [t, k]$ 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 $[a_i, g_i] \in [t, k]$. 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 $[a_i, g_i] \in [t_{i+1}, k_{i+1}]$, we have that $[a_i, g_i] = [a_i, k_{i+1}]$, and thus $[a_i, g_i] \in [t, k]$ follows from the fact that $k_{i+1}^{-1} \cdot k \in \mathbb{G}_{\alpha_i \alpha_{i+1}}$. □

Combining the above, we obtain the following by application of the reduced product construction $\mathfrak{A} \otimes \mathbb{G}$ for suitably acyclic groupoids $\mathbb{G}$, which are available according to Theorem 2.21.

Theorem 3.23. *For every $N \in \mathbb{N}$, every finite hypergraph admits a strict covering by a finite hypergraph of the from $\mathfrak{A} \otimes \mathbb{G}$ that is N-acyclic.*

As any strict hypergraph covering of any given realisation of an I-graph H induces another realisation of H by Lemma 3.20, realisations can be boosted to any desired degree of acyclicity. We have thus proved the main theorem almost as in Theorem 0.1, viz., up to the analysis of the global symmetry behaviour to which the next section will be devoted.

Theorem 3.24. *Every abstract finite specification of an overlap pattern admits, for every $N \in \mathbb{N}$, finite realisations by N-acyclic finite hypergraphs.*

4 Symmetries in reduced products

We discuss global symmetries, and especially the behaviour of our constructions under automorphisms of the input data. Section 4.1 provides the basic definitions and indicates that the matter is not entirely trivial, for two reasons.

Firstly, the straightforward, essentially relational representation of given structures or input data that we chose does not necessarily represent the relevant symmetries as automorphisms, because the representations themselves break symmetries. The most apparent example lies at the very root of our formalisations: incidence patterns I and I-graphs. If we choose a relational representation of $I = (S, E)$ as a trivial special I-graph, with singleton relations $R_e = \{(s, s')\}$ for every $e \in E[s, s']$, then the resulting structure is rigid, simply because we have individually labelled the edges by their names. Rather than this format, we need to consider I as a two-sorted structure of the form $I = (S, E, \iota_{0/1}, (\ )^{-1})$ discussed in Remark 2.2, which leaves room for non-trivial automorphisms that capture all the intended 'symmetries'.

Secondly, constructions could in principle involve choices that break symmetries of the input data or given structures. In essence, this section largely is to show that the constructions presented so far are all sufficiently natural and generic in relation to the input data, so that such choices do not occur (or at least do not have to occur). This is not trivial in all instances. In Section 4.2 we first look at those symmetries in direct and reduced products that stem from the homogeneity of the Cayley graphs of the groupoid factor in those products; Section 4.3 concerns compatibility of these constructions with symmetries that are present in the input structures. In the end we shall know that all our constructions – of realisations and coverings based on reduced products with groupoids – are indeed symmetry preserving and in themselves highly symmetric, provided the groupoids are, and that our construction of groupoids is compatible with this requirement.

In Section 4.4, we use these features of our constructions to provide a new proof of extension properties for partial automorphisms in the style of Hrushovski, Herwig and Lascar [13, 9, 11], and thus show how our generic recipe for the realisation of overlap patterns can be used to lift local symmetries, as manifested in overlaps between local substructures or local isomorphisms, to global symmetries that manifest themselves as automorphisms.

4.1 Global symmetries and automorphisms

Definition 4.1. A *symmetry* of an incidence pattern $I = (S, E)$ is an automorphism of the associated two-sorted structure, i.e., a pair[8] $\eta^I = (\eta^S, \eta^E)$ of bijections $\eta^S : S \to S$ and $\eta^E : E \to E$, such that $\eta^E(e) \in E[\eta^S(s), \eta^S(s')]$ iff $e \in E[s, s']$, and such that $\eta^E(e^{-1}) = (\eta^E(e))^{-1}$.

Definition 4.2. A *symmetry* of an I-graph $H = (V, (V_s), (R_e))$ based on an incidence pattern $I = (S, E)$ consists of a symmetry $\eta^I = (\eta^S, \eta^E)$ of I together with a permutation η^V of V such that $\eta^V(V_s) = V_{\eta^S(s)}$ and $\eta^V(R_e) := \{(\eta^V(v), \eta^V(v')) : (v, v') \in R_e\} = R_{\eta^E(e)}$.

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

Definition 4.3. A *symmetry* of an I-groupoid $\mathbb{G}$ based on an incidence pattern $I = (S, E)$ consists of a symmetry η^I of I together with a permutation η^G of G such that for all $e \in E[s, s']$, η^G maps the generator g_e of $\mathbb{G}$ to the generator $\eta^G(g_e) = g_{\eta^E(e)}$, and is compatible with the groupoid structure in the sense that, for all $s \in S$ and $g_1 \in G_{st}, g_2 \in G_{tu}$:

(i) $\eta^G(1_s) = 1_{\eta^S(s)}$;

(ii) $\eta^G(g_1 \cdot g_2) = \eta^G(g_1) \cdot \eta^G(g_2)$.

In this scenario we think of the symmetry $\eta^{\mathbb{G}}$ of the I-groupoid $\mathbb{G}$ as the triple $\eta^{\mathbb{G}} = (\eta^G, \eta^S, \eta^E)$. It follows from the definition that, dropping superscripts for notational ease, $\eta(g) \in G_{\eta(s)\eta(t)}$ for all $g \in G_{st}$ and that, for $w = e_1 \dots e_n \in E^*[s, t]$, $\eta(w^{\mathbb{G}}) = (\eta(w))^{\mathbb{G}}$ where $\eta(w) = \eta(e_1) \dots \eta(e_n) \in E^*[\eta(s), \eta(t)]$.

It is obvious that the last two definitions are compatible with the passage between I-groupoids and their Cayley graphs (cf. especially Definitions 2.13, 2.7 and 2.12).

Observation 4.4. *If $\eta^{\mathbb{G}} = (\eta^G, \eta^S, \eta^E)$ is a symmetry of the I-groupoid $\mathbb{G}$, and $H = (G, \dots)$ its Cayley graph, then $\eta^H := (\eta^V, \eta^S, \eta^E)$ for $\eta^V = \eta^G$ is also a symmetry of H. Conversely, any symmetry of an I-graph H naturally lifts to a symmetry of its completion $\bar{H}$ and of the I-groupoid $\mathbb{G} = \mathbb{G}(H)$.*

Note that the groupoid structure, i.e., the partial operation on $\mathbb{G}$ is fully determined by the Cayley graph structure (viewed as an I-graph, with E-labelled edges) together with the identification of the elements $1_s \in G_{ss}$ for $s \in S$. But, due to its homogeneity, the Cayley graph structure does not even distinguish G_{ss} within G_{*s}: it is clear that the automorphism group of the Cayley graph

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

(even as a relational structure with E-labelled edges) acts transitively on each set G_{*s}.[9]

Most importantly, the groupoid constructions from Section 2 are fully compatible with symmetries. In particular, passage from H to $\mathbb{G}(H)$ and the amalgamation constructions that lead to N-acyclic groupoids, Theorem 2.21, are such that every symmetry of I (or of the given I-graph H) lift and extend to symmetries of the resulting I-groupoid $\mathbb{G}$. The induction underlying Theorem 2.21 is based on the number of generators in the sub-groupoids, with an individual induction step according to Lemma 2.20 on the amalgamation of chains of sub-groupoids. All these notions are entirely symmetric w.r.t. any symmetries of I. We can therefore strengthen the claim of Theorem 2.21 as follows.

Corollary 4.5. *For every incidence pattern $I = (S, E)$, I-graph H and $N \in \mathbb{N}$, there are finite N-acyclic I-groupoids $\mathbb{G}$ that extend every symmetry of I and H to a symmetry of $\mathbb{G}$.*

The following notion of hypergraph automorphisms is the obvious one.

Definition 4.6. An *automorphism* of a hypergraph $\mathfrak{A} = (A, S)$ is a bijection $\eta \colon A \to A$ such that for every $s \subseteq A$: $\eta(s) := \{\eta(a) \colon a \in s\} \in S$ if, and only if, $s \in S$.

We also denote the induced bijection on S by η and may also think of an automorphism of the hypergraph $\mathfrak{A} = (A, S)$ as a pair $\eta^{\mathfrak{A}} = (\eta^A, \eta^S)$, which then automatically induces a symmetry $\eta^I = (\eta^S, \eta^E)$ of the associated intersection graph $I := I(\mathfrak{A})$ (regarded as an incidence pattern) as well as a symmetry $\eta^H = (\eta^V, \eta^S, \eta^E)$ of the I-graph representation $H(\mathfrak{A})$ of $\mathfrak{A}$. The operation of $\eta \colon A \to A$ naturally induces all the derived operations occurring in these; e.g., on the vertex set $V = \bigcup_{s \in S}(s \times \{s\})$ of $H(\mathfrak{A})$, $\eta^V(s \times \{s\}) = (\eta^S(s) \times \{\eta^S(s)\})$ where $\eta^S(s) = \{\eta(a) \colon a \in s\}$.

4.2 Groupoidal symmetries of reduced products

Consider a hypergraph $H \otimes \mathbb{G}$ or $\mathfrak{A} \otimes \mathbb{G}$ obtained as a reduced product of either an I-graph or a hypergraph with an I-groupoid (for $I = I(\mathfrak{A})$ in the hypergraph case). Any such reduced product has characteristic symmetries within its vertical fibres induced by groupoidal or Cayley symmetries of $\mathbb{G}$. In particular, these symmetries are compatible with the natural projections onto to first factor and trivial w.r.t. I.

Lemma 4.7. *The automorphism group of the reduced product $H \otimes \mathbb{G}$ between an I-graph H and an I-groupoid $\mathbb{G}$ acts transitively on the set of hyperedges*

[9] In fact, the Cayley graph of $\mathbb{G}$ consists of a disjoint union of isomorphic complete I-graphs induced on the subsets $G_{t*} = \bigcup_s G_{ts}$ for $t \in S$ (if I is connected, then these are also the connected components); the groupoid structure of $\mathbb{G}$ can be retrieved from each one of these.

$\{[V_s, g] \colon g \in G_{*s}\}$ for every $s \in S$, with trivial operation on the V_s-component. Similarly, in the direct product $H \times \mathbb{G}$, any two patches $V_s \times \{g\}$ for $g \in G_{*s}$ are bijectively related by some symmetry of $H \times \mathbb{G}$ that is trivial on the first component and on I. In the natural covering of a hypergraph $\mathfrak{A} = (A, S)$ by its reduced product $\pi \colon \mathfrak{A} \otimes \mathbb{G} \to \mathfrak{A}$ with an I-groupoid $\mathbb{G}$, where $I = I(\mathfrak{A})$, 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. All claims are immediate consequences of the homogeneity properties of the Cayley graphs of groupoids and the fact that the definitions of the direct and reduced products in question do *not* refer to the groupoid structure of $\mathbb{G}$ (with distinguished elements 1_s), other than through its Cayley graph. $\square$

4.3 Lifting structural symmetries to reduced products

Besides the vertical symmetries within fibres of coverings or realisations there is an obvious concern relating to the compatibility of reduced products with automorphisms of the input data. For these considerations, symmetries that involve non-trivial symmetries of the underlying incidence pattern I are of the essence. The aim is to show that all our constructions of realisations and coverings are sufficiently natural or canonical to allow us to lift symmetries of a hypergraph $\mathfrak{A}$ to its coverings by $\mathfrak{A} \otimes \mathbb{G}$, and of an I-graph H to its realisations obtained in direct and reduced product constructions. This requires the use of groupoids $\mathbb{G}$ that are themselves fully symmetric w.r.t. those symmetries of I that are induced by the structural symmetries of the given $\mathfrak{A}$ or H.

Recall from the discussion of Definition 4.6 that an automorphism of a hypergraph $\mathfrak{A} = (A, S)$ takes the form $\eta^{\mathfrak{A}} = (\eta^A, \eta^S)$ and canonically induces symmetries $\eta^I = (\eta^S, \eta^E)$ of $I = I(\mathfrak{A})$ and $\eta^{H(\mathfrak{A})} = (\eta^V, \eta^S, \eta^E)$ of the I-graph representation $H(\mathfrak{A})$ of $\mathfrak{A}$. If, in addition, the I-groupoid $\mathbb{G}$ admits a matching symmetry $\eta^{\mathbb{G}} = (\eta^G, \eta^I) = (\eta^G, \eta^S, \eta^E)$, based on the same $\eta^I = (\eta^S, \eta^E)$, then the covering $\pi \colon \mathfrak{A} \otimes \mathbb{G} \to \mathfrak{A}$ 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:

$$\begin{tikzcd}
\mathfrak{A} \otimes \mathbb{G} \arrow[r, "\eta^{\hat{\mathfrak{A}}}"] \arrow[d, "\pi"] & \mathfrak{A} \otimes \mathbb{G} \arrow[d, "\pi"] \\
\mathfrak{A} \arrow[r, "\eta^{\mathfrak{A}}"] & \mathfrak{A}
\end{tikzcd}$$

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

$$\eta^{\mathfrak{A} \otimes \mathbb{G}} \colon [a, g] \longmapsto [\eta^A(a), \eta^G(g)],$$

for $g \in G_{*s}$ and $a \in s$.

From Corollary 4.5 and Theorem 3.23 we thus further obtain the following.

Corollary 4.8. *Any finite hypergraph $\mathfrak{A}$ admits, for $N \in \mathbb{N}$, finite strict N-acyclic coverings $\pi \colon \hat{\mathfrak{A}} = \mathfrak{A} \otimes \mathbb{G} \to \mathfrak{A}$ by reduced products with finite N-acyclic I-groupoids $\mathbb{G}$, for $I = I(\mathfrak{A})$, such that these coverings 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 $\eta^{\hat{\mathfrak{A}}}$ of $\hat{\mathfrak{A}}$ such that $\pi \circ \eta^{\hat{\mathfrak{A}}} = \eta^{\mathfrak{A}} \circ \pi$.*

Turning to realisations of overlap patterns and their symmetries, it is equally clear that direct and reduced products $H \times \mathbb{G}$ and $H \otimes \mathbb{G}$ extend every symmetry $\eta^H = (\eta^V, \eta^S, \eta^E)$ of the I-graph $H = (V, (V_s), (R_e))$ provided the I-groupoid $\mathbb{G}$ admits a matching symmetry $\eta^{\mathbb{G}} = (\eta^G, \eta^S, \eta^E)$ with the same underlying symmetry $\eta^I = (\eta^S, \eta^E)$.

Corollary 4.9. *Every simultaneous symmetry of the I-graph H and the I-groupoid $\mathbb{G}$ gives rise to a symmetry of the direct product $H \times \mathbb{G}$ and to an automorphism of the reduced product hypergraph $H \otimes \mathbb{G}$.*

We may now combine these observations with those from Section 4.2 to obtain realisations that respect all symmetries of the overlap specification. Recall from Sections 3.2–3.4 that realisations, in the general case, were obtained in a two-stage process.

$$H \quad \longrightarrow \quad H \times \mathbb{G} \quad \longrightarrow \quad (H \times \mathbb{G}) \otimes \tilde{\mathbb{G}} \;=\; \hat{\mathfrak{A}} \quad (*)$$

The first stage takes us from the I-graph H to a direct product I-graph $H \times \mathbb{G}$ with a compatible I-groupoid $\mathbb{G}$. The second stage uses this direct product as a coherent $\tilde{I}$-graph for $\tilde{I} = I(\mathbb{G})$ (Observation 3.8), to form a reduced product with a suitable $\tilde{I}$-groupoid $\tilde{\mathbb{G}}$, which realises the overlap pattern specified in H (Corollary 3.13). The analysis of the relevant symmetries, therefore, involves structural symmetries of the input structure $H \times \mathbb{G}$ in the second stage, which also stem from groupoidal symmetries of $\mathbb{G}$.

Corollary 4.10. *For any incidence pattern $I = (S, E)$ and I-graph $H = (V, (V_s), (R_e))$, realisations $\hat{\mathfrak{A}}$ as obtained in Theorem 3.24 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 $\pi_{\hat{s}_i} \colon \hat{s}_i \to V_s$, there is a 'vertical' automorphism $\eta = \eta^{\hat{\mathfrak{A}}}$ of $\hat{\mathfrak{A}}$ that is compatible with these projections in the sense that $\pi_{\hat{s}_1} = \pi_{\hat{s}_2} \circ \eta$.*

$$\begin{tikzcd}
\hat{s}_1 \arrow[r, "\eta"] \arrow[d, "\pi_{\hat{s}_1}"'] & \hat{s}_2 \arrow[d, "\pi_{\hat{s}_2}"] \\
V_s \arrow[r, "\mathrm{id}_{V_s}"] & V_s
\end{tikzcd}$$

Proof. Consider the two-level construction indicated in $(*)$ above. It suffices to choose the two groupoids in the construction of the realisation sufficiently symmetric. For the first stage, the I-groupoid $\mathbb{G}$ can be chosen to be compatible with H and to have symmetries $\eta^{\mathbb{G}} = (\eta^G, \eta^S, \eta^E)$ based on the same $\eta^I =$

(η^S, η^E) for all symmetries $\eta^H = (\eta^V, \eta^S, \eta^E)$ of the I-graph H. Then the I-graph $H \times \mathbb{G}$ has a symmetry $\eta^{H\times\mathbb{G}}$ that simultaneously extends η^H and $\eta^{\mathbb{G}}$, for every symmetry η^H of H. In addition, $H \times \mathbb{G}$ has all the groupoidal symmetries of the Cayley graph of $\mathbb{G}$ according to Lemma 4.7.

For the second stage of the construction, $H \times \mathbb{G}$ is regarded as an $\tilde{I}$-graph where

$$\tilde{I} = I(\mathbb{G}) = (G, \tilde{E}) \quad \text{where } \tilde{E} = \bigcup_{e\in E}\big\{(g, gg_e) \colon e \in E[s,s'], g \in G_{*s}\big\}.$$

This $\tilde{I}$-graph also has the natural extension of every symmetry η^H of H as a symmetry: clearly, every symmetry of the Cayley graph of $\mathbb{G}$ that is induced by a symmetry of H extends to a symmetry of $H \times \mathbb{G}$ as an $\tilde{I}$-graph. The $\tilde{I}$-groupoid $\tilde{\mathbb{G}}$ can now be chosen compatible with $H \times \mathbb{G}$, N-acyclic for any desired level N, and such that it extends every simultaneous symmetry of $\tilde{I}$ and $H \times \mathbb{G}$. It then follows that the realisation $\hat{A} = (H \times \mathbb{G}) \otimes \tilde{\mathbb{G}}$ is compatible with every symmetry of H.

For the additional claim about 'vertical' symmetries consider two hyperedges $\hat{s}_1 = [V_s \times \{g_1\}, \tilde{g}_1]$ and $\hat{s}_2 = [V_s \times \{g_2\}, \tilde{g}_2]$ that project to the same V_s; here $g_i \in G_{*s}$ and $\tilde{g}_i \in \tilde{G}_{*g_i}$ (note that $\tilde{S} = G$ in $\tilde{I} = (\tilde{S}, \tilde{E}) = (G, \tilde{E})$).

The Cayley graph of $\mathbb{G}$ has a symmetry η_0 that takes g_1 to g_2 and whose underlying symmetry of I is trivial so that it fixes s. The corresponding simultaneous symmetry of H and the Cayley graph of $\mathbb{G}$ lifts to $H \times \mathbb{G}$ and induces matching symmetries of $\tilde{I}$ and $\tilde{\mathbb{G}}$. (This symmetry $\eta_0^{\tilde{\mathbb{G}}}$ of $\tilde{\mathbb{G}}$ will typically not be trivial with respect to $\tilde{I}$ as it links $\eta_0^{\tilde{S}}(g_1) = g_2$ in their roles as elements of $\tilde{S}$.) A purely groupoidal symmetry η_1 of $\mathbb{G}$, which is trivial w.r.t. $\tilde{I}$ and $H \times \mathbb{G}$ again, will finally suffice to align $\eta_0^{\tilde{\mathbb{G}}}(\tilde{g}_1)$ with $\tilde{g}_2$, so that the composition of η_0 and η_1 maps $\hat{s}_1 = [V_s \times \{g_1\}, \tilde{g}_1]$ to $\hat{s}_2 = [V_s \times \{g_2\}, \tilde{g}_2]$ as required. □

4.4 Lifting local to global symmetries

In its basic form, Herwig's theorem [9, 11, 10] 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 $\mathfrak{B}$.[10] It generalises a corresponding theorem by Hrushovski [13], which makes the same assertion about graphs. Elegant combinatorial proofs of both theorems can be found in the paper by Herwig and Lascar [11], which also generalises them further to classes of structures that avoid certain weak substructures (cf. Corollary 4.19 below). The variant in which all partial isomorphisms from a given collection of partial isomorphisms of the original finite structure are extended to automorphisms of one common extension is more useful for some purposes. (In general, this stronger version cannot be achieved by a straightforward iteration of the basic extension task unless the construction of $\mathfrak{B} \supseteq \mathfrak{A}$ is compatible with all automorphisms of $\mathfrak{A}$.)

[10] W.l.o.g. we may restrict attention to structures with a single relation R of some fixed arity r.

Theorem 4.11 (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, $p\colon \mathrm{dom}(p) \to \mathrm{image}(p)$ that is an isomorphism between $\mathfrak{A} \restriction \mathrm{dom}(p)$ and $\mathfrak{A} \restriction \mathrm{image}(p)$ (induced substructures). In the context of the above theorem, it is also customary to refer to the 'local symmetries', which are to be extended to global symmetries (automorphisms), as 'partial automorphisms'.

We here reproduce Herwig's theorem in an argument based on our groupoidal constructions, which may also offer a starting point for further generalisations. Before that, we prove from scratch the basic version of Herwig's theorem for a single partial isomorphism p by way of motivating our new approach to the full version below.

Excursion: the basic extension task. 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 $a \in \mathrm{dom}(p)$; we think of $\approx$ as induced by partial matchings or local bijections

$$\begin{array}{rcl} \rho_{p,n}\colon \mathrm{dom}(p) \times \{n\} & \longrightarrow & \mathrm{image}(p) \times \{n+1\} \\ (a, n) & \longmapsto & (p(a), n+1). \end{array}$$

Then, for $m \leqslant n$,

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

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

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

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 $\mathfrak{A} \times \{0\}$. In fact $\mathfrak{A}_\infty$ is a free amalgam of the slices $\mathfrak{A} \times \{n\}$.

The shift $n \mapsto n-1$ in the second component induces automorphisms $\eta\colon (a, n) \mapsto (a, n-1)$ and $\eta\colon [a, n] \mapsto [a, n-1]$ 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 $a \in \mathrm{dom}(p) \subseteq A$:

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

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

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

$$\mathfrak{B} := (\mathfrak{A} \times \mathbb{Z}_{2k})/{\approx},$$

which embeds $\mathfrak{A}$ isomorphically in the induced substructure represented by the slice $\mathfrak{A} \times \{0\}$.[11] Moreover, $\mathfrak{B}$ has the automorphism $\eta\colon [a,n] \mapsto [a,n-1]$ (now modulo $2k$ in the second component), which extends p. Therefore $\mathfrak{B}$ is a finite solution to the extension task. As $\mathfrak{B}$ is not in general compatible with (full) automorphisms of $\mathfrak{A}$, the passage from $\mathfrak{A}$ to $\mathfrak{B} \supseteq \mathfrak{A}$ that solves the extension task for one partial isomorphisms p of $\mathfrak{A}$ cannot just be iterated to further solve the extension task for another partial isomorphism p' without potentially compromising the solution for p. It turns out that suitably adapted groupoids, instead of a naive use of cyclic groups, implicitly take care of the interaction between simultaneous extension requirements.

Let us summarise this argument in light of the approach we want take below, i.e., in terms of realisations of I-graphs that specify a desired overlap pattern. For this let $I := (\{0\}, \{e_p, e_p^{-1}\})$ be the singleton incidence pattern with the two orientations of a loop for p. The overlap pattern to be realised is specified in the I-graph $H = (A, A, (p, p^{-1}))$ where A is the universe and only partition set, and p and p^{-1} stand for the partial bijections associated with the edges e_p and e_p^{-1}. For $k \in \mathbb{N}$, let

$$H \times \mathbb{Z}_k := \big(A \times \mathbb{Z}_k, (A \times \{n\})_{n \in \mathbb{Z}_k}, (\rho_{p,n}, \rho_{p,n}^{-1})_{n \in \mathbb{Z}_k}\big)$$

with partial bijections $\rho_{p,n}\colon (a,n) \mapsto (p(a), n+1)$ for $a \in \mathrm{dom}(p)$. We may think of $\mathbb{Z}_k$ as a covering $\tilde{I}$ of I in the sense of our discussion in Section 3.2; then $H \times \mathbb{Z}_k$ is the $\tilde{I}$-graph associated with the I-graph H. According to Lemma 3.7, every realisation of $H \times \mathbb{Z}_k$ induces a realisation of H. Above, we chose k such that $H \times \mathbb{Z}_k$ is coherent in the sense of Definition 2.4. Coherence of $H \times \mathbb{Z}_k$ further implies that any $\mathbb{Z}_{mk}$ for $m \geqslant 1$, viewed as an $\tilde{I}$-groupoid in the obvious manner, is compatible with $H \times \mathbb{Z}_k$. For $m \geqslant 2$, moreover, $\mathbb{Z}_{mk}$ is an m-acyclic $\tilde{I}$-groupoid. Therefore, by Observation 3.12, the reduced product $H \otimes \mathbb{Z}_{2k}$ is a realisation of $H \times \mathbb{Z}_k$, and hence also of H. Coherence, and in particular the 2-acyclicity of $\mathbb{Z}_{2k}$, implies that any relation $R^{\mathfrak{A}}$ on A for which p is a partial isomorphism lifts consistently to $H \otimes \mathbb{Z}_{2k}$, and that in this manner every slice represented by some $A \times \{n\}$ for $n \in \mathbb{Z}_{2k}$ is isomorphic to $(A, R^{\mathfrak{A}})$. The symmetry of the realisation under cyclic shifts in $\mathbb{Z}_{2k}$, finally, shows that this structure extends the partial isomorphism p of any individual slice to an automorphism. In the following we expand on this perspective.

[11] The period $\ell = 2k$ (rather than k) supports the essential condition that the domains of the partial bijections p^n and of $(p^{-1})^{\ell-n}$, which relate the same slices albeit along opposite directions in the cycle $\mathbb{Z}_\ell$, cannot be incomparable; cf. Observation 3.12 and discussion below.

A generic construction. We turn to the extension task for a specified collection P of partial isomorphisms of $\mathfrak{A} = (A, R^{\mathfrak{A}})$. By a (finite) solution to the extension task for$\mathfrak{A}$ and P we generally mean any (finite) extension $\mathfrak{B} \supseteq \mathfrak{A}$ which lifts every partial isomorphism $p \in P$ to an automorphism of $\mathfrak{B}$. With the given data $\mathfrak{A}$ and P we want to associate a natural incidence pattern I and I-graph specification of an overlap pattern, whose symmetric realisations will solve the extension task. To fit the technical constraints on overlap specifications in I-graphs according to Definition 2.3 we want to assume that P consists of strictly partial isomorphisms, i.e., that no $p \in P$ is a full automorphism of $\mathfrak{A}$. This is essentially without loss of generality, since we could add a dummy element to $\mathfrak{A}$, which is then not involved in R or in any of the $p \in P$. Our solutions to the extension task will be such that all copies of that dummy element can finally be removed from $\mathfrak{B}$ without damage.

For $I = (S, E)$ we use a 1-element set $S = \{0\}$ and endow it with one forward and one backward loop for each $p \in P$:

$$E = \{e_p \colon p \in P\} \cup \{e_p^{-1} \colon p \in P\},$$

where the e_p are pairwise distinct edge colours. For H we choose the I-graph

$$H(\mathfrak{A}, P) = (A, (R_e)_{e \in E}) \text{ where } R_e = \begin{cases} \{(a, p(a)) \colon a \in \mathrm{dom}(p)\} \text{ for } e = e_p \\ \{(p(a), a) \colon a \in \mathrm{dom}(p)\} \text{ for } e = e_p^{-1} \end{cases}$$

is the graph of p or p^{-1}, according to the orientation of e.

The idea is that the desired automorphisms will be directly induced by 'vertical shifts' in the sense of the last claim of Corollary 4.10 that link suitable pre-images of V_s in realisations of $H(\mathfrak{A}, P)$. Before we state the first version of a Herwig–Lascar theorem as Theorem 4.13 below, we introduce some terminology for useful homogeneity properties of hypergraphs.

Definition 4.12. We call a hypergraph $(B, \hat{S})$ *homogeneous* if its automorphism group acts transitively on its set of hyperedges: for $\hat{s}, \hat{s}' \in \hat{S}$ there is an automorphism η of $(B, \hat{S})$ such that $\eta(\hat{s}) = \hat{s}'$. Similarly we say that $(B, \hat{S})$ is homogeneous w.r.t. some subgroup of its full automorphisms if any two hyperedges in $\hat{S}$ are related by an automorphism from that subgroup.

Homogeneity of $(B, \hat{S})$ in particular implies that $(B, \hat{S})$ is *uniform* in the sense that all its hyperedges have the same cardinality; recall that a hypergraph is called k-*uniform* if all its hyperedges have size k.

The following variant of Herwig's theorem allows for attention to the nature of the hypergraph $(B, \hat{S})$, which describes an atlas for $\mathfrak{B}$ consisting of charts to $\mathfrak{A}$, and 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 in [12] on the basis of Herwig's theorem together with a more basic hypergraph covering result. It also stipulates a strict version of the extension of a partial isomorphism p of $\mathfrak{A}$ to an automorphism f of $\mathfrak{B} \supseteq \mathfrak{A}$ that requires not just $p \subseteq f$ but $p = f \restriction f^{-1}(A \cap f(A))$.

Theorem 4.13. *For every finite $\mathfrak{A} = (A, R^{\mathfrak{A}})$ and every set P of partial isomorphism of $\mathfrak{A}$ there is a finite relational structure $\mathfrak{B} = (B, R^{\mathfrak{B}})$ and an $|A|$-uniform hypergraph $(B, \hat{S})$ that is homogeneous w.r.t. to some subgroup of both* $\mathrm{Aut}(\mathfrak{B})$ *and* $\mathrm{Aut}(B, \hat{S})$ *such that*

(i) $\mathfrak{B} \restriction \hat{s} \simeq \mathfrak{A}$ *for every* $\hat{s} \in \hat{S}$*;*

(ii) for every $\hat{s} \in \hat{S}$, $\mathfrak{B} \subseteq \mathfrak{B} \restriction \hat{s} \simeq \mathfrak{A}$ solves the extension task for $\mathfrak{A}$ and P in the strict sense that for every $p \in P$ as applied to $\mathfrak{A} \simeq \mathfrak{B} \restriction \hat{s}$ there is some automorphism $f \in \mathrm{Aut}(\mathfrak{B})$ extending p and such that $\hat{s} \cap f(\hat{s}) = \mathrm{image}(p)$;

(iii) every non-trivial overlap between $\mathfrak{A}$-copies $\mathfrak{B} \restriction \hat{s}$ and $\mathfrak{B} \restriction \hat{t}$ as constituents of $\mathfrak{B}$ for $\hat{s}, \hat{t} \in \hat{S}$ is induced by some composition ρ_w of partial isomorphisms p and p^{-1} for $p \in P$ as applied to $\mathfrak{A} \simeq \mathfrak{B} \restriction \hat{s}$, and hence by the operation of a corresponding composition f_w of automorphisms extending the p and p^{-1} to $\mathrm{Aut}(\mathfrak{B})$ *for which* $f_w(\mathfrak{B} \restriction \hat{s}) = \mathfrak{B} \restriction \hat{t}$.

For any desired threshold $N \in \mathbb{N}$, $(\mathfrak{B}, \hat{S})$ can moreover be chosen so that the hypergraph $(B, \hat{S})$ is N-acyclic.

Proof. Assume first that P consist of strict partial isomorphisms of $\mathfrak{A}$. Let $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ be a fully symmetric and N-acyclic realisation of the overlap pattern specified in $H(\mathfrak{A}, P)$ in the sense of Corollary 4.10, with projections $\pi_{\hat{s}} \colon \hat{s} \to A$. We expand the universe $\hat{A}$ to produce an R-structure by pulling the interpretation $R^{\mathfrak{A}}$ to $\hat{A}$:

$$\mathfrak{B} := (\hat{A}, R^{\mathfrak{B}}) \quad \text{where} \quad R^{\mathfrak{B}} = \bigcup_{\hat{s} \in \hat{S}} \pi_{\hat{s}}^{-1}(R^{\mathfrak{A}}).$$

One checks that this interpretation entails that $\pi_{\hat{s}} \colon \mathfrak{B} \restriction \hat{s} \simeq \mathfrak{A}$ for every $\hat{s} \in \hat{S}$. This interpretation of R in $\mathfrak{B}$ also turns every π-compatible automorphism of the hypergraph $\hat{\mathfrak{A}}$ into an automorphism of the R-structure $\mathfrak{B}$. We want to show that it solves the extension task (in the strict sense) for every $p \in P$ and for every embedding of $\mathfrak{A}$ into $\mathfrak{B}$ via any one of the maps $\pi_{\hat{s}}^{-1}$.

By Corollary 4.10, the automorphism group of the hypergraph $\hat{\mathfrak{A}}$ acts transitively on $\hat{S}$ in a manner that is compatible with the $\pi_{\hat{s}}$. In particular, $\hat{\mathfrak{A}}$ is homogeneous w.r.t. the subgroup of those hypergraph automorphisms that are compatible with the projections and therefore also preserve the relational structure $\mathfrak{B}$. For $p \in P$, consider any hyperedge $\hat{s}$ of $\hat{\mathfrak{A}}$ and a hyperedge $\hat{s}'$ that, corresponding to condition (i) for realisations, overlaps with $\hat{s}$ as prescribed by $\rho_e = p$ (on the left-hand side of the diagram):

$$\begin{tikzcd} \hat{s} \arrow[r, "\mathrm{id}_{\hat{s}\cap\hat{s}'}"] \arrow[d, "\pi_{\hat{s}}"'] & \hat{s}' \arrow[d, "\pi_{\hat{s}'}"] \\ A \arrow[r, "p"] & A \end{tikzcd} \qquad \begin{tikzcd} \hat{s}' \arrow[r, "\eta"] \arrow[d, "\pi_{\hat{s}'}"'] & \hat{s} \arrow[d, "\pi_{\hat{s}}"] \\ A \arrow[r, "\mathrm{id}_A"] & A \end{tikzcd}$$

Consider the automorphism η of $\hat{\mathfrak{A}}$ that maps $\hat{s}'$ onto $\hat{s}$ and is compatible with $\pi_{\hat{s}'}$ and $\pi_{\hat{s}}$ (on the right-hand side of the diagram). The combination of the two diagrams shows that this automorphism η maps

$$\pi_{\hat{s}}^{-1}(\mathrm{dom}(p)) = \pi_{\hat{s}'}^{-1}(\mathrm{image}(p)) = \hat{s} \cap \hat{s}' \subseteq \hat{s}$$

onto $\pi_{\hat{s}}^{-1}(\mathrm{image}(p)) \subseteq \hat{s}$, and extends the partial map

$$p_{\hat{s}} := \pi_{\hat{s}}^{-1} \circ p \circ \pi_{\hat{s}},$$

which represents p in the embedded copy $\pi_{\hat{s}}^{-1}(\mathfrak{A}) = \mathfrak{B} \restriction \hat{s}$ of $\mathfrak{A}$. Since $\hat{s}$ and $\hat{s}'$ overlap in precisely the subsets $\pi_{\hat{s}}^{-1}(\mathrm{dom}(p)) \subseteq \hat{s}$ and $\pi_{\hat{s}'}^{-1}(\mathrm{image}(p)) \subseteq \hat{s}'$, the extension is strict in the sense of condition (ii). Condition (iii) reproduces the corresponding condition for realisations (Definition 3.5) in the current setting. N-acyclicity of the realisation $\hat{\mathfrak{A}}$ means that the hypergraph $(B, \hat{S})$ is N-acyclic.

In the case that some $p \in P$ are full automorphisms of $\mathfrak{A}$, $H(\mathfrak{A}, P)$ does not conform with Definition 2.3 of I-graphs, which requires strictly partial bijections. We therefore replace $H(\mathfrak{A}, P)$ by $H(\mathfrak{A}', P)$ for the trivial extension of $\mathfrak{A}$ by one extra dummy element $* \notin A$: $\mathfrak{A}' := \mathfrak{A} \,\dot{\cup}\, \{*\}$. If $\mathfrak{B}' \supseteq \mathfrak{A}'$ and $(B', \hat{S}')$ are obtained for $\mathfrak{A}'$ and P above, then the copy of $* \in A'$ in each $\mathfrak{B}' \restriction \hat{s}'$ can be deleted and turns the $|A'| = (|A|+1)$-uniform hypergraph $(B', \hat{S}')$ into an $|A|$-uniform hypergraph $(B, \hat{S})$ since no two hyperedges of $(B', \hat{S}')$ can overlap with their $*$-elements ($* \notin \mathrm{dom}(p) \cup \mathrm{image}(p)$, for all $p \in P$). It follows that $\mathfrak{B} = \mathfrak{B}' \restriction B$ satisfies requirements (i)–(iii). Moreover, N-acyclicity is preserved in induced substructures. □

The homogeneity and strictness properties of the solutions to the extension task obtained in our construction have further interesting consequences. In essence, we see that sufficiently acyclic realisations of $H(\mathfrak{A}, P)$ behave like 'free' solutions on a local scale, in the sense of Theorem 0.3. This is an interesting phenomenon, because really free solutions will in general be necessarily infinite. Moreover, this feature of our solutions offers a new and transparent route to the much stronger extension property of Herwig and Lascar [11] in Corollary 4.19 below.

Proposition 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 $\mathfrak{C}$ to the extension task for $\mathfrak{A}$ and P.*

The proof is essentially based on the analysis of our solutions to the extension task in the case that we use an N-acyclic realisation $(B, \hat{S})$ of the overlap specification in $H(\mathfrak{A}, P)$ (or $H(\mathfrak{A}', P) = H(\mathfrak{A} \,\dot{\cup}\, \{*\}, P)$). We draw on the alternative characterisation of hypergraph acyclicity in terms of the existence of tree decompositions, as provided by Proposition 3.4. We then prove two further claims concerning $\mathfrak{B}$ in relation to $(B, \hat{S})$ for N-acyclic $(B, \hat{S})$. Recall from Definition 3.3 that a tree decomposition of a finite hypergraph (A, S) can be represented by an enumeration of the set S of hyperedges as $S = \{s_0, \ldots, s_m\}$

such that for every $1 \leqslant \ell \leqslant m$ there is some $n(\ell) < \ell$ for which

$$s_\ell \cap \bigcup_{n<\ell} s_n \subseteq s_{n(\ell)}.$$

By Proposition 3.4, the hypergraph $(B, \hat{S})$ is N-acyclic if, and only if, every induced sub-hypergraph of size up to N admits a tree decomposition. With an N-acyclic realisation of $H(\mathfrak{A}, P)$ or $H(\mathfrak{A}', P) = H(\mathfrak{A} \dot\cup \{*\}, P))$ in our construction of a solution to the extension task for $\mathfrak{A}$, we therefore obtain the following additional properties of the hypergraph $(B, \hat{S})$ and the induced relational structure $\mathfrak{B} = (B, R^{\mathfrak{B}})$.

Claim 4.15. *In the terminology of the proof of Theorem 4.13, and for an N-acyclic realisation$(B, \hat{S})$ of $H(\mathfrak{A}, P)$ (or $H(\mathfrak{A}', P) = H(\mathfrak{A} \dot\cup \{*\}, P)$): for every substructure $\mathfrak{B}_0 \subseteq \mathfrak{B}$ of size up to N there are hyperedges $\{\hat{s}_0, \ldots, \hat{s}_m\} \subseteq \hat{S}$ such that*

(i) $(B_0, \{\hat{s}_0 \cap B_0, \ldots, \hat{s}_m \cap B_0\})$ admits a tree decomposition;

(ii) $R^{\mathfrak{B}} \restriction B_0 \subseteq \bigcup_{n \leqslant m} R^{\mathfrak{B}} \restriction \hat{s}_n$.

Proof of the claim. Let $\mathfrak{B}_0 = \mathfrak{B} \restriction B_0$, $|B_0| \leqslant N$. Due to N-acyclicity of $(B, \hat{S})$, the induced sub-hypergraph $(B_0, \{\hat{s} \cap B_0 \colon \hat{s} \in \hat{S}\})$ admits a tree decomposition represented by some enumeration of these induced hyperedges as $(\hat{s}_\ell \cap B_0)_{\ell \leqslant m}$. Due to the nature of $R^{\mathfrak{B}}$ as the union of the relations $\pi_{\hat{s}}^{-1}(R^{\mathfrak{A}})$ for $\hat{s} \in \hat{S}$, which agree in their overlaps, it is clear that condition (ii) is satisfied. $\square$

Claim 4.16. *Let $\mathfrak{B} = (B, R^{\mathfrak{B}})$ be a solution to the extension task for $\mathfrak{A}$ and P in the strict sense of Theorem 4.13, based on the hypergraph $(B, \hat{S})$. Let $\mathfrak{C} \supseteq \mathfrak{A}$ be any other (finite or infinite) solution to the extension task for $\mathfrak{A}$ and P, such that every $p \in P$ extends to an automorphism $f_p \in \mathrm{Aut}(\mathfrak{C})$.*

Let $\hat{s}, \hat{s}' \in \hat{S}$ and consider any isomorphic embedding of $\mathfrak{B} \restriction \hat{s} \simeq \mathfrak{A}$ onto an automorphic image of $\mathfrak{A}$ within $\mathfrak{C}$ of the form

$$\pi := f \circ \pi_{\hat{s}} \colon \mathfrak{B} \restriction \hat{s} \simeq f(\mathfrak{A}) \subseteq \mathfrak{C}$$

for some $f \in \mathrm{Aut}(\mathfrak{C})$. Then there is an isomorphic embedding of $\mathfrak{B} \restriction \hat{s}' \simeq \mathfrak{A}$ onto another automorphic image of $\mathfrak{A}$ within $\mathfrak{C}$ of the form:

$$\pi' := f' \circ \pi_{\hat{s}'} \colon \mathfrak{B} \restriction \hat{s}' \simeq f'(\mathfrak{A}) \subseteq \mathfrak{C}$$

for some $f' \in \mathrm{Aut}(\mathfrak{C})$ such that $\pi \cup \pi' \colon \mathfrak{B} \restriction \hat{s} \cup \mathfrak{B} \restriction \hat{s}' \to \mathfrak{C}$ is a homomorphism into $\mathfrak{C}$ from the free amalgam $\mathfrak{B} \restriction \hat{s} \cup \mathfrak{B} \restriction \hat{s}'$ of these two induced substructures over their intersection in $\mathfrak{B}$.

Proof of the claim. We observe that, according to the properties of a realisation, the overlap $\hat{s} \cap \hat{s}'$ in $(B, \hat{S})$ is induced by a path w of partial bijections p and p^{-1} for $p \in P$ in H. The corresponding composition of partial isomorphisms of $\mathfrak{A}$ gives rise to a composition of automorphisms $f_w \in \mathrm{Aut}(\mathfrak{C})$. So there is an automorphism $f' = f_w \circ f \in \mathrm{Aut}(\mathfrak{C})$ such that $f'(\mathfrak{A}) \cap f(\mathfrak{A})$ contains $\pi(\mathfrak{B} \restriction (\hat{s} \cap \hat{s}'))$. It follows that $\pi' := f' \circ \pi_{\hat{s}'}$ agrees with π on $\hat{s} \cap \hat{s}'$, whence $\pi \cup \pi'$ is well-defined on $\mathfrak{B} \restriction \hat{s} \cup \mathfrak{B} \restriction \hat{s}'$ and a homomorphism. Note that $\pi \cup \pi'$ need not be an isomorphism: it could fail to be injective due to a larger overlap between the images in $\mathfrak{C}$, and even in case it were injective, the amalgam of the images need not be free in $\mathfrak{C}$. $\square$

Proof of the proposition. Let $\mathfrak{B} = (B, R^{\mathfrak{B}})$ be a solution to the extension task for $\mathfrak{A}$ and P in the strict sense of Theorem 4.13, based on the hypergraph $(B, \hat{S})$ which is obtained as an N-acyclic realisation. Then any $\mathfrak{B}_0 \subseteq \mathfrak{B}$ of size up to N admits a tree decomposition by $(\hat{s}_0 \cap B_0)_{\ell \leqslant m}$ in the sense of Claim 4.15. By Claim 4.16, $\mathfrak{C}$ contains automorphic images $(f_\ell(\mathfrak{A}))_{\ell \leqslant m}$ of $\mathfrak{A}$ that are related to $\mathfrak{B} \restriction \hat{s}_\ell$ by individual isomorphisms of the form $\pi_\ell = f_\ell \circ \pi_{\hat{s}_\ell}$ and such that π_ℓ and $\pi_{n(\ell)}$ agree on the overlap in $\mathfrak{B}$, so that

$$\bigcup_{\ell \leqslant m} \pi_\ell(\mathfrak{B}_0 \restriction \hat{s}_\ell)$$

is a homomorphic image of $\mathfrak{B}_0$ under $\bigcup_\ell \pi_\ell$. $\square$

An analogous argument supports the following.

Remark 4.17. *For any finite R-structure $\mathfrak{A}$, any collection P of partial isomorphisms of $\mathfrak{A}$ and 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 is* N-locally free *in the sense that any substructure $\mathfrak{B}_0 \subseteq \mathfrak{B}$ of size up to N is contained in a free amalgam of copies of substructures of $\mathfrak{A}$ under* $\mathrm{Aut}(\mathfrak{B})$ *within $\mathfrak{B}$. Up to isomorphism, any such $\mathfrak{B}_0$ is thus a substructure of some free amalgam of copies of $\mathfrak{A}$.*

We also obtain the major strengthening of Theorem 4.11 due to Herwig and Lascar [11], which can be phrased as a finite-model property for the extension task within certain restricted classes of structures.

Definition 4.18. 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) solution to the extension task for $\mathfrak{A}$ and P in $\mathcal{C}$, there is a *finite* solution in $\mathcal{C}$ to this extension task.

(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 $h \colon \mathfrak{C}_i \xrightarrow{\text{hom}} \mathfrak{A}$.

The following is now immediate from Proposition 4.14.

Corollary 4.19 (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).*

Acknowledgements. I am grateful to Achim Blumensath and Julian Bitterlich for insightful comments and remarks on earlier technical reports. Thanks to Dugald McPherson and to his student Daoud Siniora for pointing out an error in an earlier version and some small corrections. Research on this topic continues to be partially supported by a grant from Deutsche Forschungsgemeinschaft on the topic of hypergraphs of controlled acyclicity (DFG grant OT 147/6-1).

References

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

[2] V. Bárány, G. Gottlob and M. Otto. Querying the guarded fragment. *Logical Methods in Computer Science*, 10 (2), pp. 1–35, 2014; also see *arXiv*:1309.5822.

[3] V. Bárány, B. ten Cate and M. Otto. Queries with guarded negation. *Proceedings of the Very Large Database Endowment*, 5(11): pp. 1328–1339, 2012; also see *arXiv*:1203.0077.

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

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

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

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

[8] E. Grädel and M. Otto. The freedoms of (guarded) bisimulation. In *Johan van Benthem on Logic and Information Dynamics*, A. Baltag, S. Smets (eds), Springer 2014, pp. 3–31.

[9] B. Herwig. Extending partial isomorphisms on finite structures. *Combinatorica*, 15, pp. 365–371, 1995.

[10] B. Herwig. Extending partial isomorphisms for the small index property of many omega-categorical structures. *Israel Journal of Mathematics*, 107, pp. 93–124, 1998.

[11] B. Herwig and D. Lascar. Extending partial isomorphisms and the profinite topology on free groups. *Transactions of the AMS*, 352, pp. 1985–2021, 2000.

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

[13] E. Hrushovski. Extending partial isomorphisms of graphs. *Combinatorica*, 12, pp. 411–416, 1992.

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

[15] D. Macpherson. A survey of homogeneous structures. *Discrete Mathematics*, 311, pp. 1599–1634, 2011.

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

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

[18] M. Otto. On groupoids and hypergraphs. Technical report, 2012. See *arXiv*:1211.5656.

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