

Amalgamation and Symmetry: From Local to Global Consistency in The Finite

Martin Otto*

Department of Mathematics
Technische Universität Darmstadt

February 2018

Abstract

Amalgamation patterns are specified by a finite collection of finite template structures together with a collection of partial isomorphisms between them. A realisation is a globally consistent solution to the locally consistent specification of this amalgamation problem. This is a single structure equipped with an atlas of distinguished substructures associated with the template structures; their overlaps realising precisely the identifications induced by the given partial isomorphisms. We present a generic construction of finite realisations of amalgamation patterns. Our construction is based on natural reduced products with suitable groupoids. The resulting realisations are generic in the sense that they can be made to preserve all symmetries of the specification. They can also be made to be universal w.r.t. to local homomorphisms up to any specified size. As key applications of the main construction we discuss finite hypergraph coverings of specified levels of acyclicity and a new route to the lifting of local symmetries to global automorphisms in finite structures.

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

*Research partially supported by DFG grant OT 147/6-1: *Constructions and Analysis in Hypergraphs of Controlled Acyclicity*. I also gratefully acknowledge participation in a programme on *Logical Structure in Computation* at the Simons Institute in Berkeley in 2016, which helped to refine some of the ideas in this work.

Contents

1	Introduction	3
2	Amalgamation patterns and their realisations	5
2.1	Amalgamation patterns	6
2.2	Measures of global consistency	10
2.3	Realisations of amalgamation patterns	12
2.4	Homomorphisms and coverings	14
2.5	Realisations from coverings	16
3	Towards generic realisations	17
3.1	An algebraic-combinatorial approach	17
3.2	Groupoidal monoids and groupoids	17
3.3	Cayley graphs of groupoids	22
3.4	Measures of global consistency in direct products	26
4	Generic realisations in reduced products	35
4.1	Degrees of acyclicity in realisations	36
4.2	Hypergraph acyclicity	36
4.3	Acyclicity in reduced products	38
4.4	Acyclicity in realisations	40
4.5	A universality property for N -acyclic realisations	41
5	Two applications	42
5.1	Hypergraph coverings	42
5.2	Lifting local to global symmetries	44

1 Introduction

At the centre of this investigation are generic algebraic-combinatorial constructions of finite amalgams of finite structures. These amalgams are based on gluing instructions specified as free amalgamations between pairs of given templates. The input data for this problem, which we call *amalgamation patterns*, consist of a finite family of finite structures as templates, equipped with a set of partial isomorphism between pairs of these. An amalgamation pattern serves as a local specification 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 designated substructure is isomorphically related to a member of the given family of templates; and the overlaps between these designated substructures is in accordance with the given family of partial isomorphisms — locally realised in pairwise disjoint amalgams. Some aspects of the combinatorial core of the matter can be illustrated in the most fundamental instance of such a problem: here the input data consists of a family of just disjoint sets together with designated partial bijections between pairs of these sets. In this scenario, a realisation is just a hypergraph: its hyperedges represent the co-ordinate domains in an atlas of charts into the given family of sets, so that changes of co-ordinates in their overlaps are induced by the given partial bijections. In the general case, we provide generic constructions of 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. The main constructions are based on reduced products with suitable finite groupoids. The existence of such finite groupoids, which need to have strong acyclicity properties in order to support global consistency, is established in [21, 20]. The relevant results from [21, 20], which generalise corresponding constructions for groups from [19], are summarised in Theorem 3.21. Our main theorem here, Theorem 4.2, states that any finite amalgamation pattern possesses highly symmetric realisations, obtained as reduced products with suitable groupoids, with additional local acyclicity and universality properties. As one key application we obtain finite branched hypergraph coverings of any desired local degree of hypergraph acyclicity. As a further application we obtain a novel 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. Our generic realisations of amalgamation patterns induced by such extension tasks (EPPA tasks in the sense of Herwig and Lascar [13]) yield not just a new route to such extensions in the style of the powerful Herwig–Lascar theorem, but apparently more generic solutions. Indeed, our formulations in Theorem 5.10 and Corollary 5.12 are more specific w.r.t. the symmetries involved, w.r.t. to the local-to-global relationship between the parts and the whole, and w.r.t. their universality properties.

At the algebraic and combinatorial level, our constructions illustrate new uses of finite groupoids and associated graph and hypergraph structures. Anal-

ogous to the established use of Cayley graphs for as a combinatorial representation of groups, we can use a Cayley graph representation of groupoids and a dual picture in a hypergraph of cosets to capture the specific acyclicity properties that are essential for our constructions. In controlling cyclic configurations in the dual hypergraph, i.e. cycles formed by overlapping cosets, the relevant acyclicity requirements exceed the scope of classical Cayley graph constructions as discussed e.g. in [2]. Classical notions of acyclicity in Cayley graphs concern the graph-theoretic girth and measure the length of generator cycles. Here we need to control *coset cycles* and measure the length of cycles formed by overlapping cosets in the Cayley graph. The technical challenge lies in the fact that the building blocks of the relevant cyclic configurations are elements of the dual hypergraph, and thus second-order objects in the Cayley graph itself, and that their size cannot a priori be bounded. Correspondingly, these acyclicity criteria are tailored to combinatorial contexts where local acyclicity for decompositions of hypergraph-like rather than graph-like structures matters. Here we use reduced products with the Cayley graphs of suitable finite groupoids to obtain natural realisations. Our first application then shows how this approach naturally lends itself to the construction of finite hypergraph coverings of controlled acyclicity. These coverings provide interesting classes of highly homogeneous and highly acyclic finite hypergraphs. As very adaptable synthetic constructions of locally acyclic hypergraphs these can play a rôle analogous to that played by known constructions of Cayley graphs of large girth in the setting of graphs. They are a source of generic examples of interest in relation to structural decomposition techniques in combinatorics and algorithmic model theory. Indeed, hypergraph acyclicity has long been recognised as an important criterion in combinatorial, algorithmic and logical contexts [6, 5]. In the well understood setting of graphs and graph-like structures, local acyclicity can be achieved in finite bisimilar unfoldings or coverings based on Cayley groups of large girth [18]. We here see that Cayley groupoids and the notion of coset acyclicity 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 [14, 19, 21, 10] w.r.t. expressive completeness results as well as w.r.t. algorithmic issues [3, 4]. The genericity of our local-to-global constructions as exemplified in their applications to extension problems for local automorphisms may have 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 [17].

At the more conceptual level, the groupoidal constructions presented here suggest interesting discrete and even finite analoga of classical concepts like branched coverings [9], notions of path independence and contractability, and the study of local symmetries [16], which invite further exploration. The algebraic-combinatorial approach to phenomena of local versus global consistency in finite relational structures may also point towards potential applications in relational models for quantum information theory as proposed in [1]. An amalgamation

pattern serves as a local specification of *pairwise* local amalgamation steps. Correspondingly, the pattern itself is indexed by an underlying graph-like structure of sites and links, the *incidence patterns* of Definition 2.5. This stands to the amalgamation task embodied in the amalgamation pattern as the intersection graph of a hypergraph stands to the actual hypergraph. And just as branched hypergraph coverings can be associated with bisimilar coverings at the level of the underlying intersection graphs [19], our realisations of amalgamation patterns involve bisimilar local unfoldings at the level of the underlying incidence patterns. Notions of bisimilarity and local unfoldings thus also play a crucial rôle in the context of local versus global symmetries and of local versus global consistency in general relational structures.

2 Amalgamation patterns and their realisations

Basic conventions. Dealing with relational structures we always appeal to a fixed finite relational signature σ and typically write $\mathcal{A} = (A, (R^A)_{R \in \sigma})$ for a σ -structure over universe A with interpretations $R^A \subseteq A^r$ for a relation symbol $R \in \sigma$ of arity $r = \text{ar}(R)$. By *substructures* $\mathcal{A}_0 \subseteq \mathcal{A}$ we always mean induced substructures $\mathcal{A}_0 = \mathcal{A} \upharpoonright A_0$ for which $R^{\mathcal{A}_0} = R^A \cap A_0^r$. This should be contrasted with *weak substructures* (otherwise often considered to be the standard notion of e.g. subgraphs) where just $R^{\mathcal{A}_0} \subseteq R^A \cap A_0^r$ is required. We shall often have occasion to consider more complex structural scenarios that involve e.g. some σ -structure together with a family of substructures, or families of partial isomorphisms within some \mathcal{A} or between the members of a family of σ -structures, et cetera. In such situations we shall often adopt a multi-sorted formalisation, with distinct sorts for different kinds of objects and an explicit encoding of relationships between objects, e.g. by means of families of functions between different sorts. As we shall see, even the labelling of such families should sometimes not be regarded as static but as a structural feature that may for instance be subject to permutations if we want to account for all relevant symmetries. The idea is illustrated in the following section for the key notion of *amalgamation patterns*, where we start from an ad-hoc preliminary formalisation in order to develop the more symmetry-aware formalisation. A *partial isomorphism* between σ -structure \mathcal{A} and \mathcal{A}' is a partial function p from A to A' that is an isomorphism between the induced substructures on its domain and image: $p: \mathcal{A} \upharpoonright \text{dom}(p) \simeq \mathcal{A}' \upharpoonright \text{image}(p)$; we write $\text{part}(\mathcal{A}, \mathcal{A}')$ for the set of all partial isomorphisms between \mathcal{A} and \mathcal{A}' .

The underlying structures that we think of as given data for amalgamation tasks will crucially be finite, but in order to discuss the intended solutions, which are also finite, we occasionally refer to related infinite structures. For many basic notions finiteness is not essential and we choose formulations that could equally be applied in infinite settings. Specific conventions will be discussed in context, but overall we just assume the standard terminology of basic universal algebra.

2.1 Amalgamation patterns

Amalgamation patterns specify intended overlaps between substructures that are modelled on templates. This specification is locally tight, and guaranteed to be locally consistent in terms of local one-to-one overlaps, but leaves unspecified the global overlap structure implicitly required for its realisations, and may not in itself be globally consistent in any straightforward sense. We give a preliminary definition to discuss first examples at a more naive level before turning to the proper definitions. Those more formal definitions of amalgamation patterns and their realisations will be slightly more involved in order to support the intended sensitivity to symmetries in the specification that we shall want to preserve in sufficiently generic constructions of realisations.

Definition 2.1. [preliminary]

An *amalgamation pattern* over some relational signature σ is a structure of type $\mathbb{H} = (\mathcal{H}, (\mathcal{A}_s)_{s \in S}, (\rho_e)_{e \in E})$ consisting of a σ -structure \mathcal{H} , whose domain $H = \dot{\bigcup}_{s \in S} \mathcal{A}_s$ is partitioned into *sites* \mathcal{A}_s of *sorts* $s \in S$, such that \mathcal{H} is the disjoint union of the induced substructures $\mathcal{A}_s = \mathcal{H} \upharpoonright \mathcal{A}_s$,

$$\mathcal{H} = \dot{\bigcup}_{s \in S} \mathcal{A}_s.$$

These substructures are related by a collection of pairwise *links* $\rho_e \in \text{part}(\mathcal{A}_s, \mathcal{A}_{s'})$, which are partial isomorphisms between \mathcal{A}_s and $\mathcal{A}_{s'}$. The index set E of links is partitioned according to $E = \dot{\bigcup}_{s, s' \in S} E[s, s']$ such that $e \in E[s, s']$ is the label of a link directed from \mathcal{A}_s to $\mathcal{A}_{s'}$.

Example 2.2. The *exploded view* of a relational structure \mathcal{A} w.r.t. a suitable family of distinguished substructures $(\mathcal{A}_s)_{s \in S}$ that cover \mathcal{A} is a special case of an amalgamation pattern. The $\mathcal{A}_s = \mathcal{A} \upharpoonright s$, for a collection of subsets $s \in S \subseteq \mathcal{P}(A)$, provide an *atlas* (in a sense to be defined below) if $\mathcal{A} = \bigcup_{s \in S} \mathcal{A}_s$. In this case, we let \mathcal{H} be the disjoint union of s -tagged copies $(\mathcal{A} \upharpoonright s) \times \{s\}$ of the \mathcal{A}_s :

$$\mathcal{H} := \dot{\bigcup}_{s \in S} (\mathcal{A} \upharpoonright s \times \{s\}).$$

Then the natural underlying indexing by a graph structure (S, E) , where $E = \{(s, s') : s \neq s', s \cap s' \neq \emptyset\}$, together with partial isomorphisms

$$\begin{aligned} \rho_e : \mathcal{A}_s \times \{s\} &\xrightarrow{\text{part}} \mathcal{A}_{s'} \times \{s'\} \\ (a, s) &\longmapsto (a, s') \text{ for } a \in s \cap s', \end{aligned}$$

record the actual identifications between elements of \mathcal{A}_s and $\mathcal{A}_{s'}$ in \mathcal{A} , for all non-trivial combinations s, s' . This very simple example holds some of the essential intuition of how an amalgamation pattern arises as an overlap specification, and also points us to what it should mean to *realise* an amalgamation pattern: clearly \mathcal{A} realises its exploded view specification. Due to its great simplicity, this class of examples trivially displays many of the special features that we shall want to guarantee through specific pre-processing in the general case.

Example 2.3. An even more special case ensues if we completely abstract away from relational content, i.e. for $\sigma = \emptyset$. Then we are dealing with the *exploded view of a hypergraph* (A, S) , which represents the hyperedges as disjoint sets together with an explicit specification of overlaps between pairs of hyperedges.

Definition 2.4. An *atlas* \mathbb{A} for a σ -structure \mathcal{A} augments \mathcal{A} according to

$$\mathbb{A} = (\mathcal{A}, U, (U_s)_{s \in S}, (\pi_{u,s})_{u \in U_s})$$

by a collection of *charts*, i.e., isomorphisms

$$\pi_{u,s}: \mathcal{A} \upharpoonright u \simeq \mathcal{A}_s$$

between induced substructures $\mathcal{A} \upharpoonright u$ and members \mathcal{A}_s of some collection of external *co-ordinate structures* $(\mathcal{A}_s)_{s \in S}$. The *co-ordinate domains* $u \subseteq A$ of the charts form a superimposed hypergraph structure (A, U) on A , with hyperedge set $U \subseteq \mathcal{P}(A)$ consisting of the domains u of the charts, which are required to cover \mathcal{A} in the sense that

$$\mathcal{A} = \bigcup_{u \in U} \mathcal{A} \upharpoonright u,$$

i.e., not just every element of A , but every tuple in the interpretation of any relation over \mathcal{A} must be fully contained in one of the hyperedges u .

We note that there may be multiple charts (into distinct, but necessarily isomorphic) co-ordinate structures \mathcal{A}_s on the same co-ordinate domain $u \subseteq A$. The collection U of co-ordinate domains thus becomes the union, but not necessarily a disjoint union, of subsets U_s of those co-ordinate domains that are associated with the co-ordinate structure \mathcal{A}_s by some $\pi_{u,s}$, for $s \in S$.

One could also associate a collection of *changes of co-ordinates* with a given atlas, which would be a collection of partial isomorphisms between co-ordinate structures \mathcal{A}_s as induced by pairs of charts $\pi_{u,s}: \mathcal{A} \upharpoonright u \simeq \mathcal{A}_s$ and $\pi_{u',s'}: \mathcal{A} \upharpoonright u' \simeq \mathcal{A}_{s'}$ with non-trivial intersection $u \cap u'$ of co-ordinate domains.

Extensive multi-sorted formalisations. In the following we shall want to avoid fixed labellings in the formalisation of complex structures like amalgamation patterns or atlases. This is essential in order to remain sensitive to all relevant symmetries, including symmetries that permute the index structure used for labelling purposes. For instance, in the harmless example of an atlas \mathbb{A} for \mathcal{A} : an automorphism of the relational structure \mathcal{A} could map each $u \in U_s$ to some $u' \in U_{s'}$ which may be matched by a permutation on the index set S and an isomorphism between the co-ordinate structures \mathcal{A}_s and $\mathcal{A}_{s'}$ that may also be compatible with the chosen co-ordinate maps $\pi_{u,s}$ and $\pi_{u',s'}$ and possibly with changes of co-ordinates (best depicted in a commuting diagram). Such a symmetry would not be apparent if the labelling of objects (such as co-ordinate domains, co-ordinate structures and co-ordinate maps and changes) were taken as static. It therefore makes sense for our concerns to adopt a multi-sorted formalisation, which allows for various functional and relational links between and within individual sorts and takes into account the sorting and typing of

various composite objects so as to trace their behaviour under the most general kinds of overall symmetries. The compressed presentation of a composite structure like an atlas as a labelled tuple of objects will still be convenient, but we shall only regard it as a shorthand for its extended multi-sorted formalisation, which in the case of an atlas would naturally involve disjoint sorts for elements of \mathcal{A} , for elements of the \mathcal{A}_s , of $U \subseteq \mathcal{P}(A)$, of a set of co-ordinate maps $P = \{\pi_{u,s} : u \in U_s, s \in S\}$, possibly also of corresponding changes of co-ordinates, and a sort for the elements of the index set S itself.

Following this idea, we firstly make explicit the underlying sites-and-links structure of an amalgamation pattern (and other related structures) by formalising it as a multigraph $\mathbb{I} = (S, E)$ according to the following definition. We shall then regard the amalgamation pattern \mathbb{H} — in the official definition to be given in Definition 2.6 below — as an *amalgamation pattern* \mathbb{H} over \mathbb{I} .

Definition 2.5. An *incidence pattern* \mathbb{I} is a finite two-sorted structure

$$\mathbb{I} = (S, E, \iota : E \rightarrow S^2, \cdot^{-1} : E \rightarrow E)$$

with sorts S (vertices, to be viewed as types of *sites* for amalgamation) and E (edges, to be viewed as types of *links*) connected by a pair of functions $\iota = (\iota_1, \iota_2) : E \rightarrow S^2$ for the allocation of source vertex $\iota_1(e) \in S$ and target vertex $\iota_2(e) \in S$ to every edge $e \in E$, and with an involutive operation of edge reversal on E , $e \mapsto e^{-1}$ subject to the requirements that $(e^{-1})^{-1} = e$ and $\iota_1(e^{-1}) = \iota_2(e)$.

The incidence pattern \mathbb{I} underlying the amalgamation pattern \mathbb{H} according to Definition 2.1 is $\mathbb{I} = (S, E, \iota, \cdot^{-1})$ where S and E are the index sets for sites and links in \mathbb{H} , with $\iota_i(e) = s_i$ determined by the requirement that $\rho_e \in \text{part}(\mathcal{A}_{s_1}, \mathcal{A}_{s_2})$ and e^{-1} determined by the condition that $\rho_{e^{-1}} = \rho_e^{-1}$. While we identify $\rho_{e^{-1}}$ with the inverse ρ_e^{-1} in the context of a given amalgamation pattern \mathbb{H} over \mathbb{I} , we keep in mind that at the level of the overlap pattern \mathbb{I} itself, the operation \cdot^{-1} is just edge reversal in a directed multigraph.

The following definition of an amalgamation pattern *over* an incidence pattern in an explicitly multi-sorted format is guided by the above considerations concerning an adequate account of symmetries.

Definition 2.6. An *amalgamation pattern* \mathbb{H} over an incidence pattern \mathbb{I} , denoted \mathbb{H}/\mathbb{I} to make the reference to \mathbb{I} explicit, is a multi-sorted structure

$$(\mathbb{I}; \mathcal{H}, \delta : H \rightarrow S; P, \eta : P \rightarrow E)$$

where

- $\mathbb{I} = (S, E, \iota, \cdot^{-1})$ is an incidence pattern (with sorts S and E as above);
- \mathcal{H} is a σ -structure whose universe H is S -partitioned by δ such that

$$\mathcal{H} = \dot{\bigcup}_{s \in S} \mathcal{A}_s$$

is a disjoint union of non-empty σ -structures $\mathcal{A}_s = \mathcal{H} \upharpoonright \delta^{-1}(s)$;

- P is a collection of partial σ -isomorphisms between pairs of these \mathcal{A}_s , E -partitioned into singleton sets (i.e. E -labelled) by η such that

$$P = \{\rho_e : e \in E\},$$

where ρ_e is the unique element of $\eta^{-1}(e)$ and such that $\rho_e \in \text{part}(\mathcal{A}_s, \mathcal{A}_{s'})$ if $\iota(e) = (s, s')$, and $\rho_{e^{-1}} = \rho_e^{-1}$.

Formally we also want to think of an atlas as a multi-sorted structure with sorts A , $\bigcup_{s \in S} \mathcal{A}_s$, $U = \bigcup_{s \in S} U_s$, $P = \{\pi_{u,s} : u \in U, s \in S\}$ and the index sort S , with the natural encoding of the labellings involved. We shall only do so for the rôle of an atlas in the context of a *realisation* of an amalgamation pattern, in Definition 2.14 below.

Symmetries. The explicitly multi-sorted formalisation of incidence and amalgamation patterns is meant to support a sufficiently broad notion of symmetries. In particular, we do want to consider symmetries of \mathbb{H}/\mathbb{I} that permute the labels $s \in S$ of sites and $e \in E$ of links in as much as such permutations respect the structure \mathbb{I} , i.e. induce a symmetry of the underlying incidence pattern. This is naturally reflected in automorphisms of the multi-sorted formalisation.

Definition 2.7. A *symmetry* of the incidence pattern $\mathbb{I} = (S, E, \iota, \cdot^{-1})$ is an automorphism of this 2-sorted structure consisting of a pair of permutations $\pi = (\pi_S, \pi_E)$ of the two sorts S and E that are compatible with ι and \cdot^{-1} .

A *symmetry* of an amalgamation pattern $\mathbb{H}/\mathbb{I} = (\mathbb{I}, \mathcal{H}, \delta, P, \eta)$ over the incidence pattern \mathbb{I} is an automorphism π of this multi-sorted structure. It is specified in terms of permutations $\pi = (\pi_S, \pi_E, \pi_H)$ of the three sorts S , E and H such that π_H is an automorphism of the relational structure \mathcal{H} , the pair (π_S, π_E) forms an automorphism of \mathbb{I} and the triple (π_S, π_E, π_H) is compatible with δ and η . A symmetry of \mathbb{H}/\mathbb{I} is *\mathbb{I} -rigid* if its operation on \mathbb{I} is trivial, i.e. if it leaves the labelling of sorts and links fixed.

In the context of \mathbb{H}/\mathbb{I} , our notational convention for compositions of partial bijections adheres to the format of multiplication/operation from the right. For links $e_1, e_2 \in E$ that match in the sense that $\iota_2(e_1) = \iota_1(e_2)$ we write $\rho_{e_1} \cdot \rho_{e_2}$ or just $\rho_{e_1} \rho_{e_2}$ for the natural (partial) composition $\rho_{e_2} \circ \rho_{e_1}$, which is a partial isomorphism from \mathcal{A}_s to $\mathcal{A}_{s'}$ for $s := \iota_1(e_1)$ and $s' := \iota_2(e_2)$. This convention naturally extends to *walks* $w = e_1 \cdots e_n$ in \mathbb{I} , where the property of a walk implies that $\iota_2(e_i) = \iota_1(e_{i+1})$ for all $i < n$, so that $\rho_w := \prod_{i=1}^n \rho_{e_i}$ stands for the composition

$$\rho_w = \prod_{i=1}^n \rho_{e_i} = \rho_{e_n} \circ \cdots \circ \rho_{e_1},$$

which is a partial isomorphism from \mathcal{A}_s to $\mathcal{A}_{s'}$ where $s = \iota_1(e_1)$ and $s' = \iota_2(e_n)$. In the same vein, we associate $\text{id}_s := \text{id}_{\mathcal{A}_s}$, the identity function on the domain \mathcal{A}_s of \mathcal{A}_s , with the unique walk λ_s of length 0 at s .

In Section 3.2 we shall appeal to the inverse semigroup structure generated by the $(\rho_e)_{e \in E}$, which may be regarded as an inverse sub-semigroup of the

symmetric inverse semigroup $I(X)$ over the set $X = H = \bigcup_{s \in S} A_s$, see e.g. [16]. Notions of *coherence*, which are to be discussed in relation to global consistency for an amalgamation pattern \mathbb{H} in the next section, can be cast in terms of this inverse semigroup $I(X)$ for $X = H$ in relation to its generators $(\rho_e)_{e \in E}$. Its elements are the compositions of the ρ_e with (partial) composition as the semigroup operation (cf. Definition 3.3). Note that all non-trivial compositions arise as compositions along walks in \mathbb{I} : products (partial compositions) can only have non-empty results if at least sites match at the interface. This is going to be converted to a groupoidal formalisation in Section 3.2. Unlike the situation in the groupoidal setting, compositions of the ρ_e in the inverse semigroup setting over \mathbb{H} are typically partial in restriction to their sites so that, for instance for $e \in E$ with $\iota(e) = (s, s')$, we need to distinguish between $\rho_{ee^{-1}} = \rho_{e^{-1}} \circ \rho_e = \rho_e^{-1} \circ \rho_e = \text{id}_{\text{dom}(\rho_e)}$ and $\rho_{\lambda_s} = \text{id}_s = \text{id}_{A_s}$.

2.2 Measures of global consistency

Definition 2.8. An amalgamation pattern \mathbb{H}/\mathbb{I} is called

- (i) *coherent* if for any two walks w_1, w_2 from s to s' , the compositions ρ_{w_1} and ρ_{w_2} agree as maps from A_s to $A_{s'}$ on the intersection of the domains.
- (ii) *simple* if for each individual link $e \in E$, $\iota(e) = (s, s')$, the partial isomorphism $\rho_e \in \text{part}(\mathcal{A}_s, \mathcal{A}_{s'})$ extends every composition ρ_w along walks w from s to s' in \mathbb{I} .
- (iii) *strongly coherent* if for any two walks w_1, w_2 from s to s' there is a walk w , also from s to s' , such that ρ_w is a common extension of ρ_{w_i} for $i = 1, 2$.

We note that coherence as defined in (i) is equivalent to the condition that any composition ρ_w along a walk w that loops at site s is a restriction of the identity $\rho_{\lambda_s} = \text{id}_{A_s}$ on A_s . It is also equivalent to the condition that the union of the bijections induced along any two walks between the same sites is again a bijection (but, unlike the case of strong coherence, not necessarily a partial isomorphisms).

Strong coherence may be viewed as a confluence property for compositions along different walks linking the same sites; it in particular implies coherence. Simplicity will be of technical interest later. It is clear that these notions impose non-trivial structural constraints on an amalgamation pattern.

Example 2.9. We may regard \mathbb{I} itself as an amalgamation pattern \mathbb{I}/\mathbb{I} (the minimal one over \mathbb{I}) in the natural manner, with S partitioned into singleton sets $\{s\}$ and singleton maps $\rho_e: \iota_1(e) \mapsto \iota_2(e)$ for $e \in E$. For any walk $w = e_1 \cdots e_n$ from $s = \iota_1(e_1)$ to $s' = \iota_2(e_n)$ in \mathbb{I} , the composition ρ_w precisely maps s to s' . This amalgamation pattern satisfies coherence, simplicity and strong coherence.

It is also instructive to check that all instances of the other most basic class of examples of amalgamation patterns, viz. exploded views according to Example 2.2 and 2.3, trivially satisfy all three conditions.

With an amalgamation pattern \mathbb{H}/\mathbb{I} we associate the equivalence relation \approx on H that is induced by the ρ_e if we regard them as identifications (in the sense

of a prescribed overlap). I.e., we let \approx be the reflexive transitive closure of the union of the graphs of the ρ_e for $e \in E$. Then for $a \in A_s$ and $a' \in A_{s'}$,

$$a \approx a' \text{ iff } a' = \rho_w(a) \text{ for some walk } w \text{ from } s \text{ to } s'.$$

In the following we write $[a]$ for the equivalence class of $a \in H$ w.r.t. \approx :

$$[a] := \{\rho_w(a) : w \text{ a walk in } \mathbb{I}, a \in \text{dom}(\rho_w)\}.$$

Lemma 2.10. *Coherence of \mathbb{H}/\mathbb{I} guarantees that \approx is trivial (coincides with equality) in restriction to each A_s . Simplicity guarantees that, for any two sites A_s and $A_{s'}$ that are directly linked by some $e \in E$ with $\iota(e) = (s, s')$, \approx identifies just the pairs linked by ρ_e . Strong coherence moreover implies that \approx is a congruence w.r.t. the interpretation of relations in \mathcal{H} .*

Proof. The first claim is an immediate consequence of the definition of coherence, Definition 2.8, and the characterisation of \approx in terms of the action of the ρ_w on H as given above. Similarly, the claim regarding simplicity is obvious from the definition. For the third claim we observe that, as a composition of partial isomorphisms ρ_e , every ρ_w is a partial isomorphism. While unions of two or more partial isomorphisms may not be partial isomorphisms (cf. Example 2.13), strong coherence in Definition 2.8 is designed to guarantee that every union of partial isomorphisms ρ_{w_i} along different walks that link the same pair A_s and $A_{s'}$ admits a common extension ρ_w . It follows that the union ρ of all those partial isomorphisms linking A_s and $A_{s'}$ is itself a partial isomorphism of this kind, say $\rho = \rho_w$. Any tuple \mathbf{a} in A_s that is (component-wise) \approx -equivalent with \mathbf{a}' in $A_{s'}$ thus is in the domain of this maximal ρ , and compatibility of $\mathbf{a} \approx \mathbf{a}'$ with the interpretations of relations R in \mathcal{A} and $A_{s'}$ follows since $\rho \in \text{part}(A_s, A_{s'})$. \square

The proof indicates that coherence deals with global consistency at the level of elements, while strong coherence enforces global consistency at the level of tuples, and therefore at the relational level. The straightforward definition of global consistency, as an outright property of a given amalgamation pattern \mathbb{H} is the following. It intuitively says that \mathbb{H} is an exploded view of a σ -structure (cf. Example 2.2).

Definition 2.11. The amalgamation pattern \mathbb{H}/\mathbb{I} is *globally consistent* if


$$\mathcal{A} := \left(\dot{\bigcup}_{s \in S} A_s \right) / \approx$$

is a well-defined σ -structure \mathcal{A} that admits an atlas with co-ordinate structures A_s and co-ordinate domains A_s / \approx with the natural charts $\pi_s : A_s / \approx \simeq A_s$ that associate the \approx -equivalence class $[a]$ of $a \in A_s$ with a .

The following is clear from Lemma 2.10.

Remark 2.12. *Global consistency implies coherence, and strong coherence implies global consistency.*

Example 2.13. Consider the following amalgamation pattern: \mathcal{A}_1 and \mathcal{A}_2 are isomorphic copies of a two-element structure consisting of a single directed edge ($\sigma = \{R\}$, R binary)



Two links identify the endpoints of these single edges in a cross-wise fashion: ρ_1 matches the source of the R -edge of \mathcal{A}_1 with the target of the R -edge of \mathcal{A}_2 , ρ_2 likewise matches the source of the R -edge of \mathcal{A}_2 with the target of the R -edge of \mathcal{A}_1 (dotted arrows in the diagram): $\rho_1: a_{1,1} \mapsto a_{2,2}$ and $\rho_2: a_{2,1} \mapsto a_{1,2}$.

Clearly, this amalgamation pattern is coherent, in fact trivially so, since the ρ_i have no non-trivial compositions; it is neither strongly coherent, nor is it globally consistent: $\rho_1 \cup \rho_2^{-1}$ is not a partial isomorphism. The twist in this overlap pattern is reminiscent of the Möbius strip: it turns out that it, too, requires an at least two-fold covering to overcome the global inconsistency in the overlap specification.

With links that are straight instead of twisted, $\rho'_1: a_{1,1} \mapsto a_{2,1}$ and $\rho'_2: a_{2,2} \mapsto a_{1,2}$, the resulting amalgamation pattern would be globally consistent (with $\mathcal{A} = (\mathcal{A}_1 \dot{\cup} \mathcal{A}_2) / \approx \simeq \mathcal{A}_i$), though still not strongly coherent, since the combination of the ρ'_i , albeit an isomorphism, is not realised as a composition of specified links.

2.3 Realisations of amalgamation patterns

The idea of a *realisation* of an amalgamation pattern strives for the converse of the passage from a structure (or hypergraph) to its exploded view. This becomes interesting in the case where the pieces in the amalgamation pattern do not already in themselves combine to form a single structure — i.e. when the local consistency built into the amalgamation pattern fails to trivially add up to the global consistency of an atlas as a global master plan according to Definition 2.11. E.g. in the example reminiscent of a Möbius strip, a two-fold covering resolves the global inconsistency.

Definition 2.14. A *realisation* of an amalgamation pattern \mathbb{H} over \mathbb{I} is a σ -structure \mathcal{A} together with an atlas

$$\mathbb{A}/\mathbb{H} = (\mathbb{H}/\mathbb{I}, \mathcal{A}, U, (U_s)_{s \in S}, (\pi_{u,s})_{u \in U_s}),$$

such that $U_s \neq \emptyset$ for all $s \in S$ and, for $u \in U_s$, $\pi_{u,s}$ is a chart from $\mathcal{A} \upharpoonright u$ onto \mathcal{A}_s , and the link structure of \mathbb{H}/\mathbb{I} is reflected in the following tight manner:

- (i) all links of \mathbb{H}/\mathbb{I} are realised as overlaps locally:
for every $u \in U_s$ and $e \in E$ with $\iota(e) = (s, s')$ there is some $u' \in U_{s'}$ such that $\text{dom}(\rho_e) = \pi_{u,s}(u \cap u')$ and $\rho_e = \pi_{u',s'} \circ \pi_{u,s}^{-1}$.

- (ii) there are no incidental overlaps, globally:
for any $u \in U_s, u' \in U_{s'}$ with $u \cap u' \neq \emptyset$, there is some walk w in \mathbb{I} such that $\text{dom}(\rho_w) = \pi_{u,s}(u \cap u')$ and $\rho_w = \pi_{u',s'} \circ \pi_{u,s}^{-1}$.

Formally we think of \mathbb{A}/\mathbb{H} as a multi-sorted structure with sorts for the elements of \mathcal{A} , the sort $U = \bigcup_{s \in S} U_s \subseteq \mathcal{P}(A)$ for the co-ordinate domains of its atlas, a sort $P = \{\pi_{u,s} : u \in U, s \in S\}$ for the charts, augmenting the multi-sorted structure \mathbb{H}/\mathbb{I} with the underlying incidence pattern \mathbb{I} , whose first sort S serves as an index sort for the atlas in the above presentation as a labelled family.

Definition 2.15. A *symmetry* of a realisation \mathbb{A}/\mathbb{H} is an automorphism of the multi-sorted structure that links the relational structure \mathcal{A} on sort A to the amalgamation pattern \mathbb{H}/\mathbb{I} through its atlas based on the sorts $U \subseteq \mathcal{P}(A)$ and $P = \{\pi_{u,s} : u \in U, s \in S\}$. Its *\mathbb{H} -rigid symmetries* are those automorphisms that fix \mathcal{H} (and thus all of \mathbb{H}/\mathbb{I} and in particular \mathbb{I}) pointwise. A realisation \mathbb{A}/\mathbb{H} of the amalgamation pattern \mathbb{H}/\mathbb{I} is *fully symmetric* over \mathbb{H}/\mathbb{I} if every symmetry of \mathbb{H}/\mathbb{I} extends to a symmetry of \mathbb{A}/\mathbb{H} , and if the automorphism group of its \mathbb{H} -rigid symmetries acts transitively on U_s , for every $s \in S$.

In terms of overlaps between isomorphic copies of the distinguished template structures \mathcal{A}_s of \mathbb{H} , condition (i) in Definition 2.14 is a richness condition. It guarantees an extension property: all specified overlaps occur wherever applicable. Condition (ii) on the other hand guarantees a minimality property: just the specified overlaps, and trivially induced ones, do occur.

Condition (i), as an *extension property*, says that (for suitable choices of $u' \in U_{s'}$ for given $u \in U_s$ and $e \in E$) the following diagram of partial maps commutes (including the information about these maps as partial isomorphisms that are bijective w.r.t. the indicated domains and ranges, which implies that the relationship between $\mathcal{A} \upharpoonright u \simeq \mathcal{A}_s$ and $\mathcal{A} \upharpoonright u' \simeq \mathcal{A}_{s'}$ is that of a disjoint amalgam):

$$\begin{array}{ccc}
& \mathcal{A} \upharpoonright (u \cap u') & \\
\pi_{u,s} \swarrow & & \searrow \pi_{u',s'} \\
\mathcal{A}_s \upharpoonright \text{dom}(\rho_e) & \xrightarrow{\rho_e} & \mathcal{A}_{s'} \upharpoonright \text{image}(\rho_e)
\end{array}$$

Similarly, condition (ii) says that (for suitable choices of $w = e_1 \cdots e_n$ for given $u \in U_s$ and $u' \in U_{s'}$) the following diagram of partial maps commutes (including the information about these maps as partial isomorphisms that are bijective w.r.t. the indicated domains and ranges):

$$\begin{array}{ccc}
& \mathcal{A} \upharpoonright (u \cap u') & \\
\pi_{u,s} \swarrow & & \searrow \pi_{u',s'} \\
\mathcal{A}_s \upharpoonright \pi_{u,s}(u \cap u') & \xrightarrow{\rho_w} & \mathcal{A}_{s'} \upharpoonright \pi_{u',s'}(u \cap u')
\end{array}$$

Given the overlaps that have to be realised according to condition (i), also any chain of compositions of ρ_{e_i} along a walk $w = e_1 \cdots e_n$ from s to s' in \mathbb{I} must

be realised as an actual sequence of overlaps that induces some intersection of the form $u \cap u'$ provided the composition ρ_w is non-empty in \mathbb{H} . Condition (ii), as a *minimality* requirement, requires that, conversely, any non-trivial intersection $u \cap u'$ arises in precisely this manner.

Condition (ii) goes beyond the requirement of global consistency, similar to the manner in which strong coherence goes beyond global consistency (cf. Definition 2.11 and discussion in Example 2.13). Strong consistency together with simplicity, however, is always sufficient to guarantee that the quotient $\mathcal{A} := (\dot{\bigcup}_{s \in S} \mathcal{A}_s) / \approx$ induces a realisation, cf. Observation 2.16 below. Also compare Remark 3.2 below for a discussion of a canonical infinite realisation for any amalgamation pattern, which can intuitively be obtained by unfolding \mathbb{H}/\mathbb{I} and \mathbb{I} itself in a tree-like fashion and gluing disjoint copies of the \mathcal{A}_s in e -related locations with overlaps as prescribed by ρ_e .

Observation 2.16. *If $(\mathcal{H}, (\mathcal{A}_s)_{s \in S}, (\rho_e)_{e \in E})$ is strongly consistent and simple, then its natural quotient w.r.t. \approx , with equivalence classes $[a] \in H / \approx$, induces a realisation \mathbb{H} / \approx based on the relational structure $\mathcal{A} = (\dot{\bigcup}_{s \in S} \mathcal{A}_s) / \approx$ and the atlas of charts with co-ordinate domains $u_s = \{[a] : a \in \mathcal{A}_s\}$ onto the $\mathcal{A}_s = \mathcal{H} \upharpoonright \mathcal{A}_s$.*

2.4 Homomorphisms and coverings

A *homomorphism*

$$\pi: \mathbb{H}^{(2)} \longrightarrow \mathbb{H}^{(1)}$$

between amalgamation patterns $\mathbb{H}^{(i)} = (\mathbb{I}^{(i)}, \mathcal{H}^{(i)}, (\mathcal{A}_s^{(i)})_{s \in S^{(i)}}, (\rho_e^{(i)})_{e \in E^{(i)}})$ over incidence patterns $\mathbb{I}^{(i)} = (S^{(i)}, E^{(i)})$, for $i = 2, 1$, consists of compatible projection maps π between the various sorts of the $\mathbb{H}^{(i)}/\mathbb{I}^{(i)}$, whose incarnation over the domain $H^{(2)}$ restricts to isomorphisms between the templates $\mathcal{A}_s^{(2)}$ of $\mathbb{H}^{(2)}$ and their images $\mathcal{A}_{\pi(s)}^{(1)}$ in $H^{(1)}$, with commuting diagrams

$$\begin{array}{ccccccc}
\mathbb{H}^{(2)}/\mathbb{I}^{(2)} & & H^{(2)} & & \mathcal{A}_s^{(2)} & \xrightarrow{\rho_e^{(2)}} & \mathcal{A}_{s'}^{(2)} & & \mathbb{I}^{(2)} & & E^{(2)} & \xrightarrow{\iota_i^{(2)}} & S^{(2)} \\
\downarrow \pi & & \downarrow \pi & & \downarrow \wr & & \downarrow \wr & & \downarrow \pi & & \downarrow & & \downarrow \\
\mathbb{H}^{(1)}/\mathbb{I}^{(1)} & & H^{(1)} & & \mathcal{A}_{\pi(s)}^{(1)} & \xrightarrow{\rho_{\pi(e)}^{(1)}} & \mathcal{A}_{\pi(s')}^{(1)} & & \mathbb{I}^{(1)} & & E^{(1)} & \xrightarrow{\iota_i^{(1)}} & S^{(1)}
\end{array}$$

Note that this notion of a homomorphism is rather strict at the local level, in that it requires local bijectivity. As the restrictions of π to the $\mathcal{A}_s^{(2)}$ are bijective, and as $\rho_{\pi(e)}^{(1)} = (\pi \upharpoonright \mathcal{A}_s^{(2)})^{-1} \circ \rho_e^{(2)} \circ \pi \upharpoonright \mathcal{A}_{s'}^{(2)}$, the compositions $\rho_e^{(2)}$ along walks of $\mathbb{I}^{(2)}$ translate directly into compositions of the $\rho_e^{(1)}$ along the image walks in $\mathbb{I}^{(1)}$, and equalities like $\rho_{w_1}^{(2)} = \rho_{w_2}^{(2)}$, $\rho_w^{(2)} = \emptyset$, or inclusions like $\rho_w^{(2)} \subseteq \text{id}_s$ at the level of $\mathbb{H}^{(2)}/\mathbb{I}^{(2)}$ imply the analogous $\rho_{\pi(w_1)}^{(1)} = \rho_{\pi(w_2)}^{(1)}$, $\rho_{\pi(w)}^{(1)} = \emptyset$, or $\rho_{\pi(w)}^{(1)} \subseteq \text{id}_{\pi(s)}$

at the level of $\mathbb{H}^{(1)}/\mathbb{I}^{(1)}$, but not vice versa. One essential way in which $\mathbb{H}^{(2)}$ can deviate from its homomorphic image in $\mathbb{H}^{(1)}$ is in terms of a potentially thinner link structure, which could involve walks in $\mathbb{I}^{(1)}$ that are not π -images of walks of $\mathbb{I}^{(2)}$, as well as in terms of a potentially richer branching in the link structure so that walks in $\mathbb{I}^{(1)}$ could arise as π -images of walks of $\mathbb{I}^{(2)}$ in multiple ways. The first of these deviations is ruled out by the following definition of a covering, which requires the homomorphism to provide lifts of walks. The second kind of variance is unaffected since lifts need not be unique. In fact, a unique lifting condition would lead to a notion of unbranched covering, while our notion crucially is one of *branched covering* (cf. Section 5.1 for a discussion of how unbranched coverings are too restrictive for our purposes).¹

Definition 2.17. A *covering* of an amalgamation pattern $\mathbb{H}^{(1)}$ over $\mathbb{I}^{(1)}$ by another amalgamation pattern $\mathbb{H}^{(2)}$ over $\mathbb{I}^{(2)}$ is a surjective homomorphism

$$\pi: \mathbb{H}^{(2)} \longrightarrow \mathbb{H}^{(1)},$$

surjective at the level of every one of the sorts involved, with the natural lifting property w.r.t. links: for every $s \in S^{(2)}$ and every $e \in E^{(1)}$ s.t. $\pi(s) = \iota^{(1)}(e)$ there exists some $e' \in E^{(2)}$ s.t. $\pi(e') = e$.

Note that the lifting property in this definition is expressed entirely in terms of the action of π on the incidence pattern, $\pi: \mathbb{I}^{(2)} \rightarrow \mathbb{I}^{(1)}$ where it stipulates that π induces a *bisimulation* (rather than just a surjective homomorphism).

Definition 2.18. A *symmetry* of a covering $\pi: \mathbb{H}^{(2)} \longrightarrow \mathbb{H}^{(1)}$ is an automorphism of the multi-sorted structure consisting of the $\mathbb{H}^{(1)}/\mathbb{I}^{(1)}$ and the maps induced by π between corresponding sorts; it is \mathbb{H} -*rigid* w.r.t. $\mathbb{H} = \mathbb{H}^{(1)}$ if it fixes the amalgamation pattern $\mathbb{H}^{(1)}/\mathbb{I}^{(1)}$ pointwise. The covering $\pi: \mathbb{H}^{(2)} \longrightarrow \mathbb{H}^{(1)}$ is *fully symmetric* (over its base $\mathbb{H}^{(1)}/\mathbb{I}^{(1)}$) if every symmetry of $\mathbb{H}^{(1)}/\mathbb{I}^{(1)}$ extends to a symmetry of the covering and if the group of its $\mathbb{H}^{(1)}$ -rigid symmetries acts transitively on each of the sets $\pi^{-1}(\{s\}) \subseteq S^{(2)}$, for $s \in S^{(1)}$.

Note that the transitivity condition for $\mathbb{H}^{(1)}$ -rigid symmetries implies that these also act transitively on each of the fibres $\pi^{-1}(a)$ above individual elements $a \in A_s^{(1)} \subseteq H^{(1)}$. The notion of a realisation in Definition 2.14 specifies required and admitted overlaps between charts in an atlas for the desired global structure just up to bisimulation. This fact is borne out in precise terms in the following.

Remark 2.19. *If $\pi: \mathbb{H}^{(2)} \longrightarrow \mathbb{H}^{(1)}$ is a covering, then any realisation $\mathbb{A}^{(2)}/\mathbb{H}^{(2)}$ induces a realisation $\mathbb{A}^{(1)}/\mathbb{H}^{(1)}$, which is obtained by the natural composition of charts of $\mathbb{A}^{(2)}$ with the local bijections induced by π .*

Proof. Let $\mathbb{A}^{(2)} = (\mathbb{H}^{(2)}/\mathbb{I}^{(2)}, \mathcal{A}, U, (U_s^{(2)})_{s \in S^{(2)}}, (\pi_{u,s}^{(2)})_{u \in U_s^{(2)}}$) be a realisation of $\mathbb{H}^{(2)}$ and put $\mathbb{A}^{(1)} = (\mathbb{H}^{(1)}/\mathbb{I}^{(1)}, \mathcal{A}, U, (U_s^{(1)})_{s \in S^{(1)}}, (\pi_{u,s}^{(1)})_{u \in U_s^{(1)}}$) where

$$\pi_{u,s}^{(1)} := \pi \circ \pi_{u,\hat{s}}^{(2)} \quad \text{for } u \in U_s^{(1)} := \bigcup \{U_{\hat{s}}^{(2)} : \pi(\hat{s}) = s\}.$$

¹E.g., the exploded view of the 3-uniform hypergraph corresponding to an apex over an n -cycle (a triangulation of an n -gon from its centre) only admits trivial unbranched coverings by disjoint copies of itself.

Clearly the new charts $\pi_{u,s}^{(1)}$ locally map $\mathcal{A} \upharpoonright u$ isomorphically onto \mathcal{A}_s for $u \in U_s^{(1)}$. One checks that conditions (i) and (ii) for realisations carry over as required. \square

2.5 Realisations from coverings

Remark 2.19 shows that the notion of a covering fits well with the intuition that an amalgamation pattern specifies the combinatorial structure of an atlas of the desired realisations just up to bisimulation. In fact, it also offers a useful route to special realisations, which will guide our further investigation.

Lemma 2.20. *For an amalgamation pattern $\mathbb{H} = (\mathbb{I}, \mathcal{H}, (\mathcal{A}_s)_{s \in S}, (\rho_e)_{e \in E})$, any covering $\pi: \hat{\mathbb{H}} \rightarrow \mathbb{H}$ by an amalgamation pattern $\hat{\mathbb{H}} = (\hat{\mathbb{I}}, \hat{\mathcal{H}}, (\mathcal{A}_{\hat{s}})_{\hat{s} \in \hat{S}}, (\rho_{\hat{e}})_{\hat{e} \in \hat{E}})$ that is simple and strongly coherent, induces a realisation of \mathbb{H} . This realisation \mathbb{A}/\mathbb{H} is based on the quotient of $\hat{\mathcal{H}}$ w.r.t. \approx , $\mathcal{A} = (\bigcup_{\hat{s} \in \hat{S}} \mathcal{A}_{\hat{s}})/\approx$.*

Proof. Let $\pi: \hat{\mathbb{H}} \rightarrow \mathbb{H}$ be a covering, $\hat{\mathbb{H}} = (\hat{H}, (\mathcal{A}_{\hat{s}})_{\hat{s} \in \hat{S}}, (\rho_{\hat{e}})_{\hat{e} \in \hat{E}})$ strongly coherent. By Remark 2.12, $\hat{\mathbb{H}}$ is in particular globally consistent so that $\mathcal{A} = \bigcup \mathcal{A}_{\hat{s}}/\approx$ is well-defined as a σ -structure on the universe $A = (\bigcup \mathcal{A}_{\hat{s}})/\approx = \{[a]: a \in \bigcup \mathcal{A}_{\hat{s}}\}$. Let $U \subseteq \mathcal{P}(A)$ consist of the subsets $u_{\hat{s}} = \{[a]: a \in \mathcal{A}_{\hat{s}}\}$ for $\hat{s} \in \hat{S}$. We put $U_s := \{u_{\hat{s}}: \pi(\hat{s}) = s\}$ (noting that we may have $u_{\hat{s}_1} = u_{\hat{s}_2}$ for $\hat{s}_1 \neq \hat{s}_2$, due to \approx -identifications). By construction \mathcal{A} admits an atlas of charts $\pi_{\hat{s}}: \mathcal{A}_{\hat{s}}/\approx \simeq \mathcal{A}_{\hat{s}}$, whose co-ordinate domains are these subsets $u_{\hat{s}} \in U$. By composition of these $\pi_{\hat{s}}$ with the covering projection π , we obtain charts

$$\pi_{u_{\hat{s}}, \pi(\hat{s})} = \pi \circ \pi_{\hat{s}}: \mathcal{A}_{\hat{s}}/\approx \simeq \mathcal{A}_{\pi(\hat{s})}$$

from co-ordinate domains $u_{\hat{s}} \in U$ onto the co-ordinate structures \mathcal{A}_s , $s = \pi(\hat{s})$, of \mathbb{H} . We check the conditions from Definition 2.14 on realisations: condition (i) uses the lifting property for the covering π , for a traversal of a single link $e \in E$, which yields a link $\hat{e} \in \hat{E}$ that is implemented as an actual, full overlap between corresponding sites $\mathcal{A}_{\hat{s}}/\approx$ and $\mathcal{A}_{\hat{s}'}/\approx$ due to simplicity of $\hat{\mathbb{H}}$. Condition (ii) for \mathbb{A}/\mathbb{H} directly corresponds to the strong coherence requirement on $\hat{\mathbb{H}}$. \square

It is instructive to see how a realisation $\mathbb{A}/\mathbb{H} = (\mathbb{H}, \mathcal{A}, U, (U_s)_{s \in S}, (\pi_{u,s})_{u \in U_s})$ of \mathbb{H} conversely also gives rise to a covering $\pi: \hat{\mathbb{H}} \rightarrow \mathbb{H}$. A natural induced covering is based on the disjoint union of u -tagged copies of $\mathcal{A} \upharpoonright u \simeq \mathcal{A}_s$ for $u \in U_s$ with some choice of $\hat{\mathbb{I}} = (\hat{S}, \hat{E})$ with $\hat{S} = \{(u, s): u \in U_s\}$. The covering link set \hat{E} can be based on a choice of pairs $\hat{e} = ((u, s), (u', s'))$ over \hat{S} , for each $e \in E$ with $\iota(e) = (s, s')$, such that (u, s) for $u \in U_s$ is paired with (u', s') for one choice of a $u' \in U_{s'}$ such that the identity on $u \cap u'$ precisely corresponds to $\pi_{u', s'} \circ \rho_e \circ \pi_{u, s}$, witnessing condition (i) for the realisation \mathbb{A}/\mathbb{H} . This covering will, however, typically fail to be strongly coherent (rather than just simple and globally consistent). This is because the project-and-lift relationship between walks in $\hat{\mathbb{I}}$ and their projections to \mathbb{I} preserves equalities between source and target sites only in the projection, not in the lifting: if \hat{w}_1 and \hat{w}_2 both link $\hat{s} = (u, s)$ to $\hat{s}' = (u', s')$ in $\hat{\mathbb{I}}$, then $\pi(\hat{w}_i) = w_i$ both link s to s' ; but some

ρ_w for which $\rho_w \supseteq \rho_{w_i}$, which exists according to condition (ii) for realisations, may lift to $\rho_{\hat{w}}$ for some \hat{w} from \hat{s} whose target site could be some other site $(u'', s') \in \pi^{-1}(s')$ rather than the desired $\hat{s}' = (u', s')$. This is in contrast with the situation of Remark 2.19, where condition (ii) for realisations is available at the level of the covering $\mathbb{H}^{(2)}$ (here $\hat{\mathbb{H}}$, above) rather than at the level of the base structure $\mathbb{H}^{(1)}$ (here \mathbb{H} , below).

3 Towards generic realisations

By *genericity* we mean to emphasise that the desired solutions to the realisation problem are fully compatible with isomorphisms between given instances. In particular our constructions do not involve any ad-hoc choices that could break internal symmetries of the given instance. Apart from greater mathematical elegance, this consideration will be of crucial importance for some of the applications to be discussed later, especially in Section 5.2.

3.1 An algebraic-combinatorial approach

The term monoid is traditionally used to relax the requirements on groups w.r.t. existence of inverses while the term groupoid, in its intended meaning here, relaxes the requirement that the fundamental binary operation be total or that the underlying carrier set be single-sorted. Being faced with the need to allow both relaxations, one could be tempted to use the term *monoidoid* for what is in fact also known as a *category*. To emphasise the salient points for our considerations, we choose to call the resulting multi-partite algebraic structures, which may not necessarily provide inverses for their partial or sort-dependent composition operation, *groupoidal monoids*. The groupoidal monoids to be considered here reflect, as algebraic structures, features of the underlying link structure \mathbb{I} of an amalgamation pattern. Towards realisations we shall later want to lift them to proper groupoid structure with inverses. Groupoids should be thought of as a multi-partite analogues of groups, as well as as categories with invertible morphisms or as close counterparts of inverse semigroups, cf. [16].

3.2 Groupoidal monoids and groupoids

Recall from Definition 2.5 the format of an incidence pattern $\mathbb{I} = (S, E)$ in multigraph notation, which stands as shorthand for the more adequate multi-sorted formalisation $\mathbb{I} = (S, E, \iota, \cdot^{-1})$. Here $\iota = (\iota_1, \iota_2)$ stands for the pair of functions that associate source and target vertices in S with every edge $e \in E$ in a manner that is compatible with edge reversal $e \mapsto e^{-1}$. A *walk* in \mathbb{I} is a finite sequence or word $w = e_1 \cdots e_n$ for some $n \in \mathbb{N}$ (the length of the walk w) such that $\iota_2(e_i) = \iota_1(e_{i+1})$ for $i < n$. We write E^* for the set of all walks in \mathbb{I} . Special walks are those of length $n = 0$: there is precisely one such at each $s \in S$, corresponding to the empty word (at s) denoted λ_s in the

following. The functions $\iota = (\iota_1, \iota_2): E \rightarrow S$ naturally extend to all of E^* , with $\iota_1(e_1 \cdots e_n) = \iota_1(e_1)$ and $\iota_2(e_1 \cdots e_n) = \iota_2(e_n)$. In particular $\iota(\lambda_s) = (s, s)$.

It is also convenient to partition E^* into the subsets of walks from s to s' in \mathbb{I} , for all pairs $s, s' \in S$,

$$E^*[s, s'] = \{w \in E^* : \iota(w) = (s, s')\},$$

and to define the partial concatenation operation on E^* as the union over

$$\begin{aligned} \cdot : E^*[s, t] \times E^*[t, u] &\longrightarrow E^*[s, u] \\ (w, w') &\longmapsto w \cdot w' := ww', \end{aligned}$$

which is precisely defined on all pairs (w, w') such that $\iota_2(w) = \iota_1(w')$.

Definition 3.1. From the incidence pattern $\mathbb{I} = (S, E, \iota, \cdot^{-1})$ we derive the *free groupoidal monoid* \mathbb{I}^* as a many-sorted structure

$$\mathbb{I}^* = (E^*, S, \iota, \cdot, \{\lambda_s : s \in S\}),$$

where E^* is the set of all walks in \mathbb{I} , with the natural extensions of ι from E to E^* as given above, with the (partial) concatenation operation \cdot for walks as above and with the (empty) walks λ_s of length 0 at s as a set of distinguished elements, which are the units w.r.t. \cdot .

Note that \mathbb{I}^* is groupoidal in that it has a sort-restricted concatenation operation, which is total in restriction to matching sorts, and that it is monoidal in not having inverses for elements other than the units.

For notational convenience we also use the shorthand

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

in close analogy with the shorthand $\mathbb{I} = (S, E)$ for $\mathbb{I} = (S, E, \iota, \cdot^{-1})$. W.r.t. these shorthands we keep in mind that they do not reflect symmetries or automorphisms appropriately in as far as they might wrongly suggest a rigid labelling of/by elements of the sets S and E .

The following idea of a *reduced product* between an amalgamation pattern \mathbb{H}/\mathbb{I} and \mathbb{I}^* yields a particularly natural, tree-like realisation of \mathbb{H}/\mathbb{I} . Albeit infinite, it illustrates and motivates key features of our construction of finite realisations. Indeed, the desired finite realisations can be pictured as suitable quotients of this canonical infinite realisation. Much of the remainder of this section focuses on the combinatorics of finite algebraic derivatives of \mathbb{I} and \mathbb{I}^* in relation to \mathbb{H}/\mathbb{I} towards the construction of reduced products, ultimately with suitable finite groupoids over \mathbb{I} in Section 4.

The *reduced product* $\mathbb{H} \otimes \mathbb{I}^*$ between \mathbb{H}/\mathbb{I} and \mathbb{I}^* is based on the quotient of the relational structure consisting of the disjoint union

$$\dot{\bigcup}_{w \in E^*} \mathcal{A}_{\iota_2(w)} \times \{w\},$$

w.r.t. to the congruence relation \approx induced as the reflexive and symmetric transitive closure of all identifications according to

$$(a, w) \approx (\rho_e(a), we)$$

for $\iota_2(w) = \iota_1(e)$. It follows that $(a_1, w_1) \approx (a_2, w_2)$ if, and only if, there are some $w_0, w'_1, w'_2 \in E^*$ such that $w_i = w_0 w'_i$ for $i = 1, 2$ and $a_2 = \rho_{w'_2}(\rho_{w'_1}^{-1}(a_1))$. Writing $[(a, w)]$ for the \approx -equivalence class of (a, w) , the natural atlas is based on the co-ordinate domains

$$u[w] := \{[(a, w)]: a \in A_s\}$$

for $s = \iota_2(w)$. In $U_s := \{u[w]: \iota_2(w) = s\}$ we then have natural charts onto \mathcal{A}_s . Choosing w_0 to be the maximal common prefix of w_1 and w_2 in the above representation $w_i = w_0 w'_i$ for $i = 1, 2$, we find the exact overlap between $u[w_1]$ and $u[w_2]$ accounted for by $\rho_{w'_2} \circ \rho_{w'_1}^{-1} = \rho_w$ for $w = w'_1{}^{-1} w'_2$. The reduced product $\mathbb{H} \otimes \mathbb{I}^* / \approx$ can be pictured as a forest of ω -branching trees of overlapping copies of the \mathcal{A}_s . The following claim is then straightforward.

Remark 3.2. *The reduced product $\mathbb{H} \otimes \mathbb{I}^* / \approx$ induces a realisation of \mathbb{H}/\mathbb{I} .*

We note that $\mathbb{H} \otimes \mathbb{I}^* / \approx$ is infinite if \mathbb{H} has at least one non-trivial link ρ_e , for which $\text{image}(\rho_e) \subsetneq A_{\iota_2(e)}$: for $a \in A_{\iota_2(e)} \setminus \text{image}(\rho_e)$ and $w_i = (e^{-1}e)^i$ (i -fold concatenation in \mathbb{I}^*), the elements (a, w_i) will be pairwise inequivalent w.r.t. \approx .

Inverse semigroup structure. To put the natural action of \mathbb{I}^* on the amalgamation pattern \mathbb{H}/\mathbb{I} on a more formal footing we may regard the inverse semigroup generated by the ρ_e as an inverse sub-semigroup of the full inverse semigroup of all partial bijections on the universe H of \mathbb{H} . The latter is known as the *symmetric inverse semigroup* on that set, in analogy with the *symmetric group* of all global bijections, cf. [16].

Definition 3.3. The *symmetric inverse semigroup* associated with a set X is the inverse semigroup $I(X) = (\text{part}(X, X), \cdot)$ consisting of all partial bijections of X with the binary operation of (partial) composition: for $\rho_i \in \text{part}(X, X)$, the product $\rho_1 \cdot \rho_2$ (typically also denoted $\rho_2 \circ \rho_1$) is the composition whose domain is $\text{dom}(\rho_1 \cdot \rho_2) = \text{dom}(\rho_1) \cap \rho_1^{-1}(\text{dom}(\rho_2)) = \rho_1^{-1}(\text{image}(\rho_1) \cap \text{dom}(\rho_2))$.

Idempotents $1_u = \text{id}_u$ (local identities) for $u \subseteq X$ serve as local left or right neutral elements w.r.t. composition with elements $\rho \in I(X)$ whose domain or range is contained in u : $1_u \cdot \rho = \rho$ if $\text{dom}(\rho) \subseteq u$ and $\rho \cdot 1_u = \rho$ if $\text{image}(\rho) \subseteq u$. And conversely, the 1_u precisely arise as compositions $\rho \cdot \rho^{-1}$ and $\rho^{-1} \cdot \rho$ for those $\rho \in I(X)$ whose domain or range is precisely u , respectively.

In the case of $I(H)$ where H is the universe of some amalgamation pattern $\mathbb{H} = (H, (\mathcal{A}_s)_{s \in S}, (\rho_e)_{e \in E})$, we shall also write id_s instead of $1_{A_s} = \text{id}_{A_s}$ for the local identities related to the partition of H into the universes A_s of the \mathcal{A}_s for $s \in S$. In this context, we may consider the partial isomorphisms ρ_e as elements of $I(H)$, the full symmetric semigroup on the set H . The inverse

semigroup composition of elements ρ_{e_1} and ρ_{e_2} is defined for any such pair, regardless of the sort-typing that is implied by the incidence pattern \mathbb{I} , or the monoidal structure of \mathbb{I}^* . For $w \in E^*$, i.e. for a walk $w = e_1 \cdots e_n$ in \mathbb{I} , the associated inverse semigroup product is

$$\rho_w = \prod_{i=1}^n \rho_{e_i} = \rho_{e_1} \cdots \rho_{e_n} = \rho_{e_n} \circ \cdots \circ \rho_{e_1}$$

in the sense of partial composition of maps, in agreement with the monoidal product of \mathbb{I}^* . In this case the ι -values encode source and target sites that accordingly determine inclusions of the form $\text{dom}(\rho_w) \subseteq A_s$ for $s = \iota_1(w)$ and $\text{image}(\rho_w) \subseteq A_{s'}$ for $s' = \iota_2(w)$. But the inverse semigroup structure of $I(H)$ allows arbitrary products $\prod_{i=1}^n \rho_{e_i}$ regardless of ι -values. On one hand it partly reflects the above inclusion constraints in equations like $\rho_w = \text{id}_s \cdot \rho_w$ and $\rho_w = \rho_w \cdot \text{id}_{s'}$; on the other hand, this reflection is just partial because all such equations trivialise for $\rho_w = \emptyset$ (which may happen for products along walks in \mathbb{I} that are good in the sense of \mathbb{I}^* , and always happens for products that are not aligned along walks in \mathbb{I} and hence not represented in \mathbb{I}^* at all). In particular $\prod_{i=1}^n \rho_{e_i} = \emptyset$ for $e_1 \cdots e_n \notin E^*$.

Definition 3.4. The inverse semigroup $I(\mathbb{H}) \subseteq I(H)$ is generated as a sub-semigroup of $I(H)$ by the $(\rho_e)_{e \in E}$ of $\mathbb{H} = (H, (\mathcal{A}_s)_{s \in S}, (\rho_e)_{e \in E})$ together with the local identities $\text{id}_s := \text{id}_{A_s}$.²

Since we regard the local identities id_s as induced by walks λ_s of length 0 at $s \in S$, at least all the *non-empty* elements are partial isomorphisms of the form $\rho_w = \prod_{i=1}^n \rho_{e_i}$ for walks w in \mathbb{I} , i.e. for $w \in E^*$. If we restrict to just those elements of $I(\mathbb{H})$ that are induced by compositions ρ_w of ρ_e along walks w in \mathbb{I} , we move away from the globally defined composition operation of an inverse semigroup.

Groupoid structure. We now want to look at groupoids whose partial and sort-sensitive operation of full rather than partial composition reflects the incidence pattern \mathbb{I} as the underlying combinatorial schema of sorts. Groupoids can also be described as categories all of whose morphisms isomorphisms. Fixing an underlying incidence pattern \mathbb{I} with sorts S and E means that we fix the object sort of the category to be the set S (of sites or sorts) and a set of generators $e \in E$ for the morphism sort that link the objects of sort s and s' as prescribed by $\iota(e) = (s, s')$. In these terms, the following definition puts the focus on the algebraic composition structure for the morphisms.

Definition 3.5. A *groupoid* \mathbb{G} over the incidence pattern \mathbb{I} , denoted by the shorthand \mathbb{G}/\mathbb{I} , is a multi-sorted structure of the form

$$\mathbb{G}/\mathbb{I} = (\mathbb{I}, G, \iota: G \rightarrow S^2, \cdot)$$

²Note that local identities id_s will not necessarily arise as non-trivial products of ρ_e with their inverses, since all ρ_e with $\iota_1(e) = s$ may have $\text{dom}(\rho_e) \subsetneq A_s$.

whose set G of groupoid elements is partitioned by ι into the sets $G[s, s'] := \iota^{-1}(s, s') = \{g \in G : \iota(g) = (s, s')\}$ of groupoid elements of source/target sorts s/s' , with a partial composition operation defined on pairs of matching interface sort according to

$$\begin{aligned} \cdot : G[s, s'] \times G[s', s''] &\longrightarrow G[s, s''] \\ (g_1, g_2) &\longmapsto g_1 \cdot g_2 =: g_1 g_2 \end{aligned}$$

that is associative and has, for all $s, s' \in S$,

- (i) unique left and right neutral elements $1_s \in G[s, s]$, s.t.
 $1_s g = g$ for all $g \in G[s, s']$ and $g 1_s = g$ for all $g \in G[s', s]$;
- (ii) unique inverses $g^{-1} \in G[s', s]$ for $g \in G[s, s']$, s.t.
 $g g^{-1} = 1_s$ and $g^{-1} g = 1_{s'}$.

In addition to these basic algebraic conditions we require G to be generated by E (the link sort of \mathbb{I}), via a map $e \mapsto e^{\mathbb{G}}$ that associates a groupoid element $e^{\mathbb{G}} \in G[\iota_1(e), \iota_2(e)]$ with every $e \in E$, such that $(e^{-1})^{\mathbb{G}} = (e^{\mathbb{G}})^{-1}$, and such that

- (iii) every $g \in G$ can be written as a product $g = \prod_{i=1}^n e_i^{\mathbb{G}} =: w^{\mathbb{G}}$ for some walk $w = e_1 \cdots e_n \in E^*$ in \mathbb{I} .⁽³⁾

It is convenient to use the enumeration of relevant ingredients as shorthand:

$$\mathbb{G} = (G, (G[s, s'])_{s, s' \in S}, \cdot, (1_s)_{s \in S}, (e^{\mathbb{G}})_{e \in E}).$$

We also use the suggestive abbreviations

$$G[* , s] := \bigcup_{s' \in S} G[s', s] \quad \text{and} \quad G[s, *] := \bigcup_{s' \in S} G[s, s'],$$

so that, for instance, \cdot is defined precisely on $\bigcup_{s \in S} (G[* , s] \times G[s, *])$.

Definition 3.6. \mathbb{G}/\mathbb{I} is called *simple* if the generators $(e^{\mathbb{G}})_{e \in E}$ are pairwise distinct and $e^{\mathbb{G}} \neq 1_s$ for all $e \in E, s \in S$.⁴

For a simple groupoid \mathbb{G}/\mathbb{I} we identify e with $e^{\mathbb{G}}$ and regard the link sort E of \mathbb{I} as a subset of G .

Example 3.7. Looking at \mathbb{I}/\mathbb{I} as an amalgamation pattern as in Example 2.9, and at the action of \mathbb{I}^* on this particular amalgamation pattern, we can isolate a natural groupoid $\mathbb{G}(\mathbb{I})/\mathbb{I}$ within the semigroup $I(\mathbb{I}) \subseteq I(S)$: it consists of those $\rho_w \in I(\mathbb{I})$ that are induced by walks $w \in \mathbb{I}^*$, with $\iota(\rho_w) = \iota(w)$ and with the natural composition. Since these ρ_w uniformly are just the singleton maps $\rho_w : \iota_1(w) \mapsto \iota_2(w)$, their composition in matching interface sorts is exact rather than partial, and induces a groupoid operation with units $1_s = \text{id}_{\{s\}} = \rho_{\lambda_s}$. This groupoid $\mathbb{G}(\mathbb{I})/\mathbb{I}$ is simple if, and only if $\mathbb{I} = (S, E)$ is a simple graph (rather than a multi-graph, possibly with loops).

³The length 0 walk λ_s at s is taken to generate the (empty) product 1_s of sort $\iota(1_s) = (s, s)$. Note that $w^{\mathbb{G}} \in G[s, s']$ if w is a walk from s to s' in \mathbb{I} , i.e. $\iota(w^{\mathbb{G}}) = \iota(w)$.

⁴ $e^{\mathbb{G}} \neq 1_s$ is meaningful as a constraint just for reflexive links $e \in E[s, s]$ (loops at s).

Example 3.8. Call an amalgamation pattern \mathbb{H}/\mathbb{I} *complete* if all the ρ_e are bijections between the sites involved: $\rho_e: A_{\iota_1(e)} \rightarrow A_{\iota_2(e)}$ is a bijection for all $e \in E$. In this case, too, the groupoidal action of \mathbb{I}^* induces a natural groupoid $\mathbb{G}(\mathbb{H})/\mathbb{I}$ within the semigroup $I(\mathbb{H}) \subseteq I(H)$. Again, $\mathbb{G}(\mathbb{H})$ consists of the $\rho_w \in I(H)$ induced by walks $w \in \mathbb{I}^*$, with the natural composition so that, due to completeness, every ρ_w is a bijection between $\text{dom}(\rho_w) = A_{\iota_1(w)}$ and $\text{image}(\rho_w) = A_{\iota_2(w)}$.

Remark 3.9. *The free groupoidal monoid \mathbb{I}^* induces a simple groupoid obtained as the quotient of \mathbb{I}^* w.r.t. cancellation of all factors ee^{-1} , so that ww^{-1} is identified with $\lambda_{\iota_1(w)}$, and w^{-1} inverse to w . We regard this groupoid, which is a simple groupoid over \mathbb{I} and infinite for non-trivial \mathbb{I} , as the free groupoid over \mathbb{I} . The same groupoid up to isomorphism is obtained in the manner of Example 3.7 from the natural tree unfolding of the multi-graph \mathbb{I} .*

Towards the construction of realisations we shall be interested in groupoids over \mathbb{I} that respect those algebraic identities that are induced by the inverse semigroup action of \mathbb{I}^* on a given amalgamation pattern over \mathbb{I} , as expressed in $I(\mathbb{H})$ (cf. Definition 3.4). The corresponding notion of *compatibility* is defined as follows.

Definition 3.10. A groupoid \mathbb{G}/\mathbb{I} is *compatible* with the amalgamation pattern \mathbb{H}/\mathbb{I} , both over the same incidence pattern \mathbb{I} , if, for all $w \in \mathbb{I}^*$,

$$w^{\mathbb{G}} = 1_s \text{ (in } \mathbb{G}) \Rightarrow \rho_w \subseteq \text{id}_s = \text{id}_{A_s} \text{ (in } I(\mathbb{H})).$$

Lemma 3.11. *For any amalgamation pattern \mathbb{H}/\mathbb{I} , the following conditions are equivalent*

- (i) \mathbb{H}/\mathbb{I} is coherent;
- (ii) every groupoid \mathbb{G}/\mathbb{I} over the same \mathbb{I} is compatible with \mathbb{H}/\mathbb{I} ;
- (iii) the groupoid $\mathbb{G}(\mathbb{I})/\mathbb{I}$ from Example 3.7 is compatible with \mathbb{H}/\mathbb{I} .

Proof. By Definition 2.8, \mathbb{H}/\mathbb{I} is coherent if any ρ_w induced by a walk in \mathbb{I} that loops at site s satisfies $\rho_w \subseteq \text{id}_s$ in \mathbb{H} . If $w^{\mathbb{G}} = 1_s$ in any \mathbb{G}/\mathbb{I} , then $\iota(w) = (s, s)$ implies that w is a walk that loops at s , whence coherence implies that $\rho_w \subseteq \text{id}_s$ as required for compatibility. Conversely, compatibility of the groupoid $\mathbb{G}(\mathbb{I})/\mathbb{I}$ with \mathbb{H}/\mathbb{I} implies that \mathbb{H}/\mathbb{I} is coherent: in this groupoid, $w^{\mathbb{G}(\mathbb{I})} = 1_s$ precisely for those walks w in \mathbb{I} with $\iota(w) = (s, s)$; so compatibility implies that $\rho_w \subseteq \text{id}_s$ in \mathbb{H} for all those, as required for coherence. \square

3.3 Cayley graphs of groupoids

With a groupoid \mathbb{G}/\mathbb{I} we associate a *Cayley graph*. First and foremost, this is the natural generalisation to the setting of groupoids of the notion of a Cayley graph for groups w.r.t. to a given set of generators. Moreover, it provides us with a crucial link between amalgamation patterns \mathbb{H}/\mathbb{I} and groupoids \mathbb{G}/\mathbb{I} : the Cayley graph of \mathbb{G} is an amalgamation pattern, indeed a very special, homogeneous amalgamation pattern over the same underlying incidence pattern \mathbb{I} .

A group G with generator set $E \subseteq G$ that is closed under inverses, can be seen as a special case of a groupoid, based on an incidence pattern $\mathbb{I} = (S, E)$ with a singleton set $S = \{0\}$ of sites, and reflexive links $e \in E$ with trivial $\iota(e) = (0, 0)$. This format of \mathbb{I} implies that \cdot is total and indeed (G, \cdot) a group; conversely, any group with generator set $E = E^{-1}$ can be cast in this format. In this sense, the following definition subsumes the familiar notion of the Cayley graph of a group (or of a Cayley group and its associated Cayley graph) as a special case.

Definition 3.12. With a groupoid \mathbb{G}/\mathbb{I} over the incidence pattern \mathbb{I} associate its *Cayley graph*, which is the many-sorted (multi-)graph

$$\mathbb{G}/\mathbb{I} = (\mathbb{I}, G, \delta: G \rightarrow S, R, \eta: R \rightarrow E),$$

with vertex set $G = \bigcup_{s \in S} G[* , s]$, vertex-labelled by S w.r.t. this partition, which is formalised by the partition function δ , and with the E -labelled edge set

$$R = \bigcup_{e \in E} R_e \quad \text{with} \quad R_e = \{(g, ge^{\mathbb{G}}) : g \in G[* , \iota_1(e)]\},$$

whose labelling is formalised by the partition map η .

Observation 3.13. *The Cayley graph of \mathbb{G}/\mathbb{I} has the format of an amalgamation pattern, where ρ_e is the partial bijection induced by right multiplication with the generator $e^{\mathbb{G}}$, whose graph is R_e . Note that ρ_e bijectively relates the associated sites $G[* , s]$ and $G[* , s']$ if $\iota(e) = (s, s')$. This amalgamation pattern is complete in the sense of Example 3.8. Indeed the groupoid \mathbb{G} can be retrieved as the groupoid induced by this amalgamation pattern in the manner discussed in that example.*

The simplified, labelled representation of the underlying amalgamation pattern is

$$(G, (G[* , s])_{s \in S}, (\rho_e)_{e \in E}).$$

Note that inversion corresponds to edge reversal, $R_{e^{-1}} = R_e^{-1}$, and that walks in \mathbb{I} represent groupoid multiplication by groupoid elements as represented as products of generators.

In the following the notation \mathbb{G}/\mathbb{I} can refer to \mathbb{G}/\mathbb{I} in these interchangeable contexts: the groupoid, its Cayley graph or the associated amalgamation pattern. Differences arise especially w.r.t. symmetries, as will be discussed below.

Symmetries. Similar to Definition 2.7 for amalgamation patterns, we also use the term *symmetry* in connection with groupoids and their Cayley graphs to refer to the adequate notion of automorphisms of the multi-sorted structures that allows for the parallel permutation of all the sorts involved, rather than fixing, say, the labelling of sites, links, or generators. We recall that it is largely to this end that we have, e.g. officially made the underlying incidence pattern \mathbb{I} an internal component of the structure of a groupoid *over* \mathbb{I} .

Definition 3.14. A *symmetry* of a groupoid $\mathbb{G}/\mathbb{I} = (\mathbb{I}, G, \iota, \cdot)$ over the incidence pattern $\mathbb{I} = (S, E, \iota, \cdot^{-1})$ is an automorphism π of this multi-sorted structure. It is specified in terms of permutations $\pi = (\pi_S, \pi_E, \pi_G)$ of the three sorts S , E and G that bijectively map each sort within itself such that the pair (π_S, π_E) is an automorphism of \mathbb{I} and the triple (π_S, π_E, π_G) is compatible with $\iota^{\mathbb{G}} = (\iota_1, \iota_2): G \rightarrow S^2$ and with the groupoid operation \cdot on G .

The groupoid \mathbb{G}/\mathbb{I} is *fully symmetric over \mathbb{I}* if every symmetry of \mathbb{I} extends to a symmetry of \mathbb{G}/\mathbb{I} .

Symmetries of the Cayley graph of \mathbb{G}/\mathbb{I} are defined in agreement with the notion of symmetries of amalgamation patterns in Definition 2.7. Correspondingly, a symmetry of the Cayley graph of \mathbb{G}/\mathbb{I} is called *\mathbb{I} -rigid* if it fixes \mathbb{I} pointwise.

We may think of a symmetry of the groupoid \mathbb{G}/\mathbb{I} either as an extension of an automorphism of the two-sorted structure \mathbb{I} , or as an automorphism of the algebraic groupoid structure \mathbb{G} that also permutes generators in a manner compatible with the structure of \mathbb{I} , i.e. inducing an automorphism of \mathbb{I} .

Despite the close relationship, there are crucial differences between the groupoid on one hand and its Cayley graph or amalgamation pattern on the other. These differences are particularly apparent in relation to their symmetries. In particular, we see that the groupoid operation \cdot cannot be definable *within* the Cayley graph, although \mathbb{G} is isomorphic to the groupoid generated from the groupoidal action of \mathbb{I}^* on this Cayley graph according to Example 3.8 — which means that the abstract algebraic structure of the groupoid is fully determined by the Cayley graph. Definability of the groupoid operation within the Cayley graph is ruled out by richer symmetries of the latter.⁵ In fact, the group of automorphism of the Cayley graph acts transitively on each partition set $G[* , s]$: any pair g, g' from the same partition set $G[* , s]$ are related by an automorphism that is induced by left multiplication with $g'g^{-1}$ on $G[\iota_1(g), *]$, and trivial elsewhere:

$$\begin{aligned} G &\longrightarrow G \\ h &\longmapsto \begin{cases} g'g^{-1}h & \text{for } h \in G[\iota_1(g), *] \\ h & \text{else.} \end{cases} \end{aligned}$$

This bijectively maps $G[\iota_1(g), s']$ onto $G[\iota_1(g'), s']$, for every s' , and in particular maps g to g' . Since it preserves both the vertex partition and the edge relations R_e , it is an automorphism of the Cayley graph, in fact an \mathbb{I} -rigid one.

Remark 3.15. *The group of \mathbb{I} -rigid symmetries of the Cayley graph of \mathbb{G}/\mathbb{I} acts transitively on the subsets $G[* , s]$ for every $s \in S$.*

The \mathbb{I} -rigid automorphism group of the groupoid, on the other hand, is trivial, since it preserves each $G[s', s]$ as a set and in particular fixes $1_s \in G[s, s]$ and all generator elements $e^{\mathbb{G}}$ for $e \in E$. All automorphisms of the Cayley groupoid, not just the \mathbb{I} -rigid ones, must preserve the set $\{1_s : s \in S\}$ of units.

⁵The relationship is similar to that between an affine space and its associated vector space: the latter is definable *from* the former, as the linear space of translations; it fails to be definable *within* the former because there is no canonical choice of an origin.

Measures of acyclicity. As an amalgamation pattern, the Cayley graph of a groupoid \mathbb{G}/\mathbb{I} need in general neither be simple nor coherent, let alone strongly coherent. Due to completeness, coherence and simplicity fall into one; and both conditions precisely mean that the groupoid is degenerate in the sense that $G[s, s] = \{1_s\}$ for all s . Corresponding criteria for groupoids and products of amalgamation patterns with groupoids will be studied below. These considerations are closely related to the following notion of qualified *coset acyclicity* for groupoids, which generalises corresponding notions for Cayley groups from [18]. Coset acyclicity puts much more severe restrictions than the familiar notion of *large girth*: while girth concerns the length of cycles formed by generators (i.e. of graph-theoretic cycles in the Cayley graph), coset acyclicity concerns the length of cycles formed by overlapping cosets w.r.t. subgroup(oid)s generated by subsets of the set of generators (i.e. of hypergraph-theoretic cycles in a dual hypergraph associated with the Cayley graph).

Definition 3.16. In a groupoid \mathbb{G}/\mathbb{I} , a subset $\alpha \subseteq E$ that is closed under inversion ($\alpha = \alpha^{-1}$) induces a *sub-groupoid* $\mathbb{G}[\alpha] \subseteq \mathbb{G}$ generated by α . $\mathbb{G}[\alpha]/\mathbb{I}$ consists of groupoid elements in $G[\alpha] := \{w^{\mathbb{G}} : w \in \mathbb{I}^* \upharpoonright \alpha\} \subseteq G$, with groupoid operation induced by \mathbb{G} . We write $G[\alpha, s, s']$ for the set of elements $g \in G[\alpha]$ with $\iota(g) = (s, s')$, which is the set of those $g \in G[s, s']$ that admit a representation as a product of generators in α , including empty products of the form $\lambda_s^{\mathbb{G}} = 1_s$ for all $s \in S$.⁶ The (left) α -coset at $g \in G[* , s]$ is

$$g\mathbb{G}[\alpha] = \{gh : h \in \mathbb{G}[\alpha], \iota_1(h) = s\} = \{gh : h \in \mathbb{G}[\alpha, s, *]\}.$$

Definition 3.17. For $n \geq 2$, a *coset cycle* of length n in the groupoid \mathbb{G}/\mathbb{I} , is a cyclically indexed tuple of *pointed cosets* $(g_i, g_i\mathbb{G}[\alpha_i])_{i \in \mathbb{Z}_n}$ such that, for all i :

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

The first condition in this definition of a coset cycle says that consecutive cosets are linked in the sense that they overlap in the named representatives; the second condition says that there is no direct shortcut within any one of these cosets, from its immediate predecessor to its immediate successor.

Definition 3.18. For $N \geq 2$, a Cayley groupoid \mathbb{G}/\mathbb{I} is called *N-acyclic* if it does not have any coset cycles of lengths $2 \leq n \leq N$.

The weakest meaningful coset acyclicity condition, 2-acyclicity, deserves special attention. If we think of α -cosets as α -connected components, then the following observation suggest to picture 2-acyclicity as a condition of simple connectivity: the intersection of two connected pieces is itself connected.

Observation 3.19. For any groupoid \mathbb{G}/\mathbb{I} : \mathbb{G} is 2-acyclic if, and only if, for all $\alpha_1, \alpha_2 \subseteq E$ such that $\alpha_i = \alpha_i^{-1}$,

$$\mathbb{G}[\alpha_0] \cap \mathbb{G}[\alpha_1] = \mathbb{G}[\alpha_0 \cap \alpha_1].$$

⁶Due to the trivial nature of the extra components, $\mathbb{G}[\alpha]$ could also be cast as a groupoid over the restricted incidence pattern $\mathbb{I} \upharpoonright \alpha$ over $S \upharpoonright \alpha = \iota_1(\alpha) = \iota_2(\alpha)$.

Proof. Generally, $\mathbb{G}[\alpha_0 \cap \alpha_1] \subseteq \mathbb{G}[\alpha_0] \cap \mathbb{G}[\alpha_1]$.

A coset cycle of length 2 consists of two pointed cosets, $g_0\mathbb{G}[\alpha_0]$ at g_0 and $g_1\mathbb{G}[\alpha_1]$ at g_1 , such that $g_0, g_1 \in g_0\mathbb{G}[\alpha_0] \cap g_1\mathbb{G}[\alpha_1]$ while $g_0\mathbb{G}[\alpha_0 \cap \alpha_1]$ is disjoint from $g_1\mathbb{G}[\alpha_0 \cap \alpha_1]$. Therefore

$$g_0g_1^{-1} \in (\mathbb{G}[\alpha_0] \cap \mathbb{G}[\alpha_1]) \setminus \mathbb{G}[\alpha_0 \cap \alpha_1]$$

shows that $\mathbb{G}[\alpha_0] \cap \mathbb{G}[\alpha_1] \neq \mathbb{G}[\alpha_0 \cap \alpha_1]$. Conversely, any $h \in (\mathbb{G}[\alpha_0] \cap \mathbb{G}[\alpha_1]) \setminus \mathbb{G}[\alpha_0 \cap \alpha_1]$ gives rise to a coset cycle consisting of $\mathbb{G}[\alpha_0]$ at 1_s for $s = \iota_1(h)$, and $h\mathbb{G}[\alpha_1]$ at h . \square

Remark 3.20. *The free groupoid over \mathbb{I} of Remark 3.9, which is infinite in non-trivial cases, is acyclic in the sense that it has no coset cycles.*

The following is shown in [21, 20], in a construction that considerably extends the corresponding, much simpler result for Cayley groups from [19].

Theorem 3.21. *For every finite incidence pattern \mathbb{I} and $N \geq 2$, there are finite groupoids \mathbb{G}/\mathbb{I} that are N -acyclic. Moreover, such \mathbb{G}/\mathbb{I} can be chosen to be simple and compatible with any given amalgamation pattern \mathbb{H}/\mathbb{I} , as well as compatible with any symmetries of the underlying amalgamation pattern \mathbb{H}/\mathbb{I} in the sense that every symmetry of \mathbb{H}/\mathbb{I} extends to a symmetry of \mathbb{G}/\mathbb{I} .*

Compare Definition 3.10 in connection with the compatibility assertion, and Definitions 3.14 and 2.7 in connection with symmetries of \mathbb{G}/\mathbb{I} and \mathbb{H}/\mathbb{I} , respectively. We isolate the symmetry preservation properties in the following notion of being *fully symmetric*. The above then says that finite N -acyclic groupoids can be obtained that are fully symmetric over the given data.

Recall from Definition 3.14 that a groupoid \mathbb{G}/\mathbb{I} is *fully symmetric* over \mathbb{I} if every symmetry of the incidence pattern \mathbb{I} extends to a symmetry of the groupoid \mathbb{G} .

Definition 3.22. A groupoid \mathbb{G}/\mathbb{I} is called *fully symmetric over \mathbb{H}/\mathbb{I}* if every symmetry of \mathbb{I} that is induced by a symmetry of the amalgamation pattern \mathbb{H}/\mathbb{I} extends to a symmetry of the groupoid \mathbb{G} .

Note that the definition only refers to symmetries of the groupoid \mathbb{G}/\mathbb{I} ; regarding the rich additional structure of \mathbb{I} -rigid automorphisms of the Cayley graph of \mathbb{G}/\mathbb{I} over and above these, see Remark 3.15.

3.4 Measures of global consistency in direct products

Direct products of amalgamation patterns \mathbb{H}/\mathbb{I} with suitable groupoids will allow us to obtain strongly coherent and simple coverings $\hat{\mathbb{H}}/\hat{\mathbb{I}}$, and, through these by Lemma 2.20, realisations. The ensuing realisations themselves thus materialise as reduced products, viz. quotients of the coverings $\hat{\mathbb{H}}/\hat{\mathbb{I}}$, which are obtained as direct products, w.r.t. the congruence relation \approx induced by the identifications according to the ρ_ε of $\hat{\mathbb{H}}$.

The qualitative difference between coherence and strong coherence may be interpreted as a fundamental difference between element-wise versus tuple-wise consistency.⁷ At the technical level this difference manifests itself as a difference between graph-like and hypergraph-like phenomena. While graph-like structures admit (even unbranched) finite N -acyclic coverings through direct products with Cayley graphs of groups [18, 8], finite N -acyclic hypergraph coverings (in general necessarily branched coverings) naturally arise as reduced products with Cayley graphs of groupoids.

We show that direct products of amalgamation patterns with Cayley graphs of compatible finite groupoids \mathbb{G}/\mathbb{I} produce finite coverings that are coherent, while further direct products with Cayley graphs of compatible and 2-acyclic groupoids guarantee simplicity and strong coherence. The coverings involved really arise at the level of the underlying incidence pattern \mathbb{I} , where they can be pictured as *bisimilar coverings* in the usual graph setting (where an edge- and vertex-coloured graph encodes a Kripke structure).

In this sense also a groupoid \mathbb{G}/\mathbb{I} induces a covering $\hat{\mathbb{I}}$ of \mathbb{I} by the Cayley graph of \mathbb{G} , as follows. In a groupoid \mathbb{G}/\mathbb{I} , the projection map $\pi: g \mapsto \iota_2(g)$ induces an unbranched bisimilar covering of the multi-graph structure \mathbb{I} by the multi-graph structure of the Cayley graph of \mathbb{G} . By this we mean a homomorphism, mapping edges $(g, ge^{\mathbb{G}})$ at $g \in G$ with $\iota_2(g) = \iota_1(e)$ to the edge e in \mathbb{I} , which also satisfies this unique lifting property: at every $g \in G$ with $\iota_2(g) = \iota_1(e)$, e uniquely lifts to $(g, ge^{\mathbb{G}})$. Putting

$$\hat{E} := \{e[g] = (g, ge^{\mathbb{G}}) : e \in E, g \in G, \iota_2(g) = \iota_1(e)\},$$

we obtain an incidence pattern $\hat{\mathbb{I}}$, which is based on G as its sort of sites and \hat{E} as its sort of links:

$$\hat{\mathbb{I}} = \hat{\mathbb{I}}(\mathbb{G}) = (G, \hat{E}, \iota, \cdot^{-1}).$$

This covering $\pi: \hat{\mathbb{I}} \rightarrow \mathbb{I}$ is a rendering of the Cayley graph of the groupoid \mathbb{G}/\mathbb{I} — refined to become an incidence pattern with individually labelled edges in $\hat{E} = \bigcup_{e \in E} Re$. With the following definition we lift this natural covering relationship between $\hat{\mathbb{I}}(\mathbb{G})$ and \mathbb{I} to a covering $\pi: \hat{\mathbb{H}} \rightarrow \mathbb{H}$ between amalgamation patterns $\hat{\mathbb{H}} = \hat{\mathbb{H}} \otimes \mathbb{G}$ over $\hat{\mathbb{I}}$ and \mathbb{H} over \mathbb{I} .

Definition 3.23. The *direct product* of an amalgamation pattern

$$\mathbb{H}/\mathbb{I} = (\mathbb{I}, \mathcal{H}, \delta: H \rightarrow S, P, \eta: P \rightarrow E)$$

over the incidence pattern $\mathbb{I} = (S, E, \iota, \cdot^{-1})$ with a groupoid \mathbb{G}/\mathbb{I} is the amalgamation pattern

$$\mathbb{H} \otimes \mathbb{G} := (\hat{\mathbb{I}}, \mathcal{H} \otimes \mathbb{G}, \hat{\delta}: H \otimes G \rightarrow G, P \otimes G, \hat{\eta}: P \otimes G \rightarrow \hat{E})$$

⁷There seems to be a vague analogy between these discrete combinatorial notions and classical geometric notions of path independence and contractibility in homotopy.

over the incidence pattern $\hat{\mathbb{H}} = (G, \hat{E}, \iota, \cdot^{-1})$ induced by \mathbb{G}/\mathbb{I} . The universe $H \otimes G$ of $\mathcal{H} \otimes \mathbb{G}$ is the set

$$\begin{aligned} H \otimes G &= \{(a, g) \in H \times G : \delta(a) = \iota_2(g)\} \\ &= \{(a, g) \in H \times G : a \in A_s, g \in G[* , s] \text{ for some } s \in S\}, \end{aligned}$$

partitioned by $\hat{\delta}$ which is the projection to the second component. $\mathcal{H} \otimes \mathbb{G}$ as a σ -structure is the disjoint union of the $\mathcal{A}_g := \mathcal{A}_s \times \{g\}$ for $g \in G[* , s], s \in S$; $P \otimes G$ correspondingly consists of all pairs $(p, g) \in P \times G$ where $\iota_1(\eta(p)) = \iota_2(g)$ regarded as a partial isomorphism between \mathcal{A}_g and $\mathcal{A}_{ge^{\mathbb{G}}}$ in the natural manner.

Compare Definitions 2.17 and 3.22 for the following.

Lemma 3.24. *The natural projection from the direct product $\mathbb{H} \otimes \mathbb{I}$ to the first factor \mathbb{H} , $\pi_1 : (a, g) \mapsto a$, is a covering of the amalgamation pattern \mathbb{H}/\mathbb{I} by $\mathbb{H} \otimes \mathbb{G}$ as an amalgamation pattern over the incidence pattern $\hat{\mathbb{H}}$ induced by the Cayley graph of \mathbb{G} . If \mathbb{G}/\mathbb{I} is fully symmetric over \mathbb{H}/\mathbb{I} , then this covering is fully symmetric over \mathbb{H}/\mathbb{I} .*

Proof. That π_1 satisfies the conditions for a covering is straightforward. For the lifting property in particular, any ρ_e at the level of \mathbb{H} lifts to $\rho_e \times g : (a, g) \mapsto (\rho_e(a), ge^{\mathbb{G}})$ at every $g \in \mathbb{G}[* , \iota_2(g)]$ for $\iota_2(g) = \iota_1(e)$; this is the partial isomorphism between \mathcal{A}_g and $\mathcal{A}_{ge^{\mathbb{G}}}$ associated in $\mathbb{H} \otimes \mathbb{G}$ with the link $e[g] = (g, ge^{\mathbb{G}}) = (g, ge^{\mathbb{G}}) \in \hat{E}$. We note that this further implies that the following diagram, in which the vertical isomorphisms are induced by π_1 , commutes for any pair of walks \hat{w} in $\hat{\mathbb{H}}$ and w in \mathbb{I} that are in such a projection/lifting relationship, and where therefore $g' = gw^{\mathbb{G}}$:

$$\begin{array}{ccc} \mathcal{A}_g & \xrightarrow{\rho_{\hat{w}}} & \mathcal{A}_{g'} \\ \downarrow \cong & & \downarrow \cong \\ \mathcal{A}_{\iota_2(g)} & \xrightarrow{\rho_w} & \mathcal{A}_{\iota_2(g')} \end{array}$$

For the symmetry claim we first observe that any pair of symmetries of the Cayley graph of \mathbb{G}/\mathbb{I} and of \mathbb{H}/\mathbb{I} that agree w.r.t. their action on \mathbb{I} extends to a symmetry of $\mathbb{H} \otimes \mathbb{G}$ and indeed of the covering relationship between $\mathbb{H} \otimes \mathbb{G}$ and \mathbb{H} . The salient point here is that $\mathbb{H} \otimes \mathbb{G}$ is defined in terms of the Cayley graph of \mathbb{G}/\mathbb{I} (as an amalgamation pattern) rather than in terms of \mathbb{G} as a groupoid. It is therefore clear that, if \mathbb{G}/\mathbb{I} is fully symmetric over \mathbb{H}/\mathbb{I} , every symmetry of \mathbb{H}/\mathbb{I} extends to a symmetry of the covering. For the claim of full symmetry of $\mathbb{H} \otimes \mathbb{G}$ over \mathbb{H} , in this situation, it remains to argue that the group of \mathbb{H} -rigid symmetries acts transitively on each of the sets $\pi^{-1}(\{s\}) = G[* , s] \subseteq G$ (regarded as sets of sites of $\hat{\mathbb{H}}$), which is immediate from Remark 3.15: the \mathbb{I} -rigid automorphisms of the Cayley graph of \mathbb{G} act transitively on the sets $G[* , s]$. \square

Lemma 3.25. *For the covering of \mathbb{H}/\mathbb{I} by $\mathbb{H} \otimes \mathbb{G}/\hat{\mathbb{H}}$:*

- (i) if \mathbb{G}/\mathbb{I} is compatible with \mathbb{H}/\mathbb{I} , then $\mathbb{H} \otimes \mathbb{G}/\hat{\mathbb{I}}$ is coherent;
(ii) if \mathbb{H}/\mathbb{I} is coherent and if \mathbb{G}/\mathbb{I} is simple and 2-acyclic,⁸ then $\mathbb{H} \otimes \mathbb{G}/\hat{\mathbb{I}}$ is simple and strongly coherent.

Proof. In (i), compatibility of \mathbb{G}/\mathbb{I} with \mathbb{H}/\mathbb{I} means that, for every $w \in \mathbb{I}^*$ with $w^{\mathbb{G}} = 1_s$ in \mathbb{G} , we must have $\rho_w \subseteq \text{id}_{A_s} = \text{id}_s$ in \mathbb{H} , where ρ_w is the composition of the ρ_e along the walk from s to s described by w in \mathbb{I} . Coherence of $\mathbb{H} \otimes \mathbb{G}$ on the other hand requires that for every walk \hat{w} in $\hat{\mathbb{I}}$ that loops back from some g to the same g , $\rho_{\hat{w}} \subseteq \text{id}_{A_g}$, where $\rho_{\hat{w}}$ is the composition of the ρ_e along the walk \hat{w} . But that \hat{w} loops at g in $\hat{\mathbb{I}}$, means that $w^{\mathbb{G}} = 1_s$ for the natural projection w of \hat{w} induced by $\pi: e[g] = (g, ge^{\mathbb{G}}) \mapsto e$ and for $s = \iota_1(w)$. This projection also associates the action of $\hat{\mathbb{I}}^*$ on $\hat{\mathbb{H}} = \mathbb{H} \otimes \mathbb{G}/\hat{\mathbb{I}}$ with the action of \mathbb{I}^* on $\mathbb{H} = \mathbb{H}/\mathbb{I}$, as indicated in the commuting diagram above, so that

$$\rho_{\hat{w}}(a, g) = (\rho_w(a), gw^{\mathbb{G}}).$$

So, if $w^{\mathbb{G}} = 1_s$ and therefore $\rho_w \subseteq \text{id}_{A_s}$ by compatibility of \mathbb{G} , it follows that $\rho_{\hat{w}} \subseteq \text{id}_{A_g}$ as required for coherence of $\hat{\mathbb{H}}$.

For (ii), consider now a direct product $\hat{\mathbb{H}} = \mathbb{H} \otimes \mathbb{G}$ where \mathbb{G}/\mathbb{I} is both compatible with \mathbb{H}/\mathbb{I} and 2-acyclic and assume that \mathbb{H}/\mathbb{I} is coherent. Coherence of \mathbb{H}/\mathbb{I} means that the equivalence relation \approx , which is induced by identification of linked elements $a \approx \rho_e(a)$ (for a, e with $a \in \text{dom}(\rho_e)$) never identifies any two elements in the same A_s (cf. Lemma 2.10). In other words, \approx -equivalence classes $[a]$ of elements $a \in H$ intersect any one partition set $A_s \subseteq H$ in at most one element. So there is, for every $a \in H$, a partial function $f_a: S \rightarrow H$ such that $[a] \cap A_s = \{f_a(s)\}$ if $f_a(s)$ is defined, and empty otherwise.⁹ In particular, every $a \in H$ defines a subset $\alpha_a \subseteq E$ of generators according to

$$\alpha_a := \{e \in E: f_a(s) \in \text{dom}(\rho_e) \text{ for } s = \iota_1(e)\},$$

which can be seen as the set of links that carry a . The set α_a is closed under the converse operation of \mathbb{I} , or closed under inverses as a generator set in \mathbb{G} , since $f_a(s) \in \text{dom}(\rho_e)$ (and $f_a(s)$ defined) for $s = \iota_1(e)$ if, and only if, $f_a(s') \in \text{dom}(\rho_{e^{-1}})$ for $s' = \iota_1(e^{-1}) = \iota_2(s)$ (and $f_a(s')$ defined).

Towards strong coherence of $\hat{\mathbb{H}}/\hat{\mathbb{I}}$, consider walks \hat{w}_i in $\hat{\mathbb{I}}$ that link the same sites g and $g' = gh = gw_i^{\mathbb{G}}$, where the w_i are the projections of \hat{w}_i to \mathbb{I} , for $i = 1, 2$. The partial isomorphisms $\rho_{\hat{w}_i}$ between \mathcal{A}_g and $\mathcal{A}_{g'}$ are lifts of their projections ρ_{w_i} between \mathcal{A}_s and $\mathcal{A}_{s'}$, for $s = \iota_2(g) = \iota_1(w_i)$ and $s' = \iota_2(g') = \iota_2(w_i)$. Let $d_i = \text{dom}(\rho_{w_i}) \subseteq A_s$, and put

$$\alpha_i := \bigcap_{a \in d_i} \alpha_a,$$

which is the set of those links of \mathbb{H} that preserve the \approx -equivalent copies of all elements of d_i . By 2-acyclicity of \mathbb{G} , the pointed cosets $(g, g\mathbb{G}[\alpha_1])$ and

⁸Note that \mathbb{G}/\mathbb{I} is compatible with \mathbb{H}/\mathbb{I} by Lemma 3.11.

⁹Think of f_a as describing a partial section that traces the element a through the desired realisation.

$(g', g' \mathbb{G}[\alpha_2])$ cannot form a coset cycle. Since ρ_{w_1} maps d_1 onto its image at g' , w_1 is composed of generators in α_1 , whence $g' \in g \mathbb{G}[\alpha_1]$; similarly, $g \in g' \mathbb{G}[\alpha_2]$ since ρ_{w_2} maps d_2 onto its image at g' (or its inverse maps this image onto d_2). So the second condition of Definition 3.17 must be violated, which for this potential 2-cycle means that g, g' are in the same coset w.r.t. $\alpha := \alpha_1 \cap \alpha_2$, i.e. that

$$w_1^{\mathbb{G}} = w_2^{\mathbb{G}} = w^{\mathbb{G}} \text{ for some } w \in \mathbb{G}[\alpha].$$

It follows that $d_1 \cup d_2 \subseteq \text{dom}(\rho_w)$. For the unique lift of this walk w in \mathbb{I} to a walk \hat{w} from g to g' therefore

$$\rho_{\hat{w}} \supseteq \rho_{\hat{w}_i}$$

provides a common extension. This shows strong coherence of $\hat{\mathbb{H}}/\hat{\mathbb{I}}$.

Moreover, simplicity and 2-acyclicity of \mathbb{G} and strong coherence of $\hat{\mathbb{H}}/\hat{\mathbb{I}}$ together imply simplicity of $\hat{\mathbb{H}}/\hat{\mathbb{I}}$, as shown by a simple indirect argument. (We write $e^{\mathbb{G}} = e$ in this setting of a simple \mathbb{G} .) By strong coherence there is a maximal element among all the $\rho_{\hat{w}}$ that link \mathcal{A}_g and \mathcal{A}_{ge} (note that $ge \neq g$ since \mathbb{G} is simple, which implies that $\hat{\mathbb{I}}$ is simple in that \hat{E} has no loops (or multiple edges)). If this $\rho_{\hat{w}}$ were not the same as $\rho_{(g, ge)}$ as required for simplicity, then the domain of $\rho_{\hat{w}}$ would be strictly larger than $\text{dom}(\rho_{(g, ge)})$. In projection to \mathbb{H} , $\text{dom}(\rho_e) \subsetneq \text{dom}(\rho_w)$ implies that w cannot use the generator e , i.e. $w \in \mathbb{G}[\alpha]$ for some $\alpha \subseteq E \setminus \{e, e^{-1}\}$. But then $(g, g \mathbb{G}[\alpha])$ and $(ge, ge \mathbb{G}[e, e^{-1}])$ would form a coset cycle of length 2 in \mathbb{G} . \square

For the following compare Definition 2.18 w.r.t. symmetries of coverings.

Corollary 3.26. *Any finite amalgamation pattern \mathbb{H}/\mathbb{I} possesses a covering by a finite, simple and strongly coherent amalgamation pattern $\tilde{\mathbb{H}}/\tilde{\mathbb{I}}$. Moreover, the covering $\pi: \tilde{\mathbb{H}}/\tilde{\mathbb{I}} \rightarrow \mathbb{H}/\mathbb{I}$ can be chosen to be fully symmetric over \mathbb{H}/\mathbb{I} .*

Proof. For $\tilde{\mathbb{H}}$ we take a two-fold direct product $(\mathbb{H} \otimes \mathbb{G}) \otimes \hat{\mathbb{G}}$ as obtained in Lemma 3.25, which is a covering of (a covering of) \mathbb{H}/\mathbb{I} by Lemma 3.24, and simple and strongly coherent for suitable choices of \mathbb{G} and $\hat{\mathbb{G}}$, by the previous lemma. We may choose \mathbb{G} compatible with \mathbb{H}/\mathbb{I} , fully symmetric over \mathbb{H}/\mathbb{I} , and $\hat{\mathbb{G}}$ simple and fully symmetric over $(\mathbb{H} \otimes \mathbb{G})/\hat{\mathbb{I}}$, so that $\tilde{\mathbb{H}}/\tilde{\mathbb{I}}$ is fully symmetric over \mathbb{H}/\mathbb{I} : any symmetry of \mathbb{H}/\mathbb{I} first lifts to $(\mathbb{H} \otimes \mathbb{G})/\hat{\mathbb{I}}$ and then further to $\tilde{\mathbb{H}}/\tilde{\mathbb{I}}$. The required \mathbb{H} -rigid symmetries that relate sites of the covering over the same \mathcal{A}_s of \mathbb{H}/\mathbb{I} are induced by lifts of \mathbb{I} -rigid symmetries of the Cayley graph of \mathbb{G} to symmetries of $(\mathbb{H} \otimes \mathbb{G}) \otimes \hat{\mathbb{G}}$ (recall that the group of \mathbb{I} -rigid symmetries of the Cayley graph of \mathbb{G} is transitive on the sets $G[* , s]$ by Remark 3.15). Any such symmetry extends to an \mathbb{H} -rigid symmetry of $\mathbb{H} \otimes \mathbb{G}$: \mathbb{I} -rigid symmetries of the Cayley graph of \mathbb{G} induce \mathbb{H} -rigid symmetries of $\mathbb{H} \otimes \mathbb{G}$ by Lemma 3.24. Since $\hat{\mathbb{G}}$ is fully symmetric over $(\mathbb{H} \otimes \mathbb{G})/\hat{\mathbb{I}}$, it extends this symmetry of $\mathbb{H} \otimes \mathbb{G}$ further to a symmetry of $(\mathbb{H} \otimes \mathbb{G}) \otimes \hat{\mathbb{G}}$, which remains \mathbb{H} -rigid, again by Lemma 3.24. \square

The following analysis shows that a covering as in Corollary 3.26 can even be obtained as a direct product $\mathbb{H} \otimes \tilde{\mathbb{G}}$ for a suitable finite groupoid $\tilde{\mathbb{G}}/\mathbb{I}$ in one

step. In other words, we want to replace the twofold direct product construction, which first guarantees coherence, then strong coherence at the level of a partially unfolded incidence pattern ($\tilde{\mathbb{I}}$ as an unbranched covering of \mathbb{I}), by a single direct product.

Let us consider an iterated direct product $\tilde{\mathbb{H}} := (\mathbb{H} \otimes \mathbb{G}) \otimes \hat{\mathbb{G}}$ as from the above construction, i.e. based on the following:

- \mathbb{H}/\mathbb{I} and \mathbb{G}/\mathbb{I} are over $\mathbb{I} = (S, E)$, \mathbb{G} is simple;
- \mathbb{G} is compatible with \mathbb{H} (so that $\mathbb{H} \otimes \mathbb{G}$ is coherent by Lemma 3.25);
- $(\mathbb{H} \otimes \mathbb{G})/\hat{\mathbb{I}} = \hat{\mathbb{H}}/\hat{\mathbb{I}}$ and $\hat{\mathbb{G}}$ are over $\hat{\mathbb{I}} = (G, \hat{E})$, the incidence pattern induced by the Cayley graph of \mathbb{G} , as an unbranched bisimilar covering $\pi: \hat{\mathbb{I}} \rightarrow \mathbb{I}$, where $\pi: g \mapsto \iota_2(g)$;
- $\hat{\mathbb{G}}/\hat{\mathbb{I}}$ is simple and 2-acyclic (and by Lemma 3.11 compatible with $\hat{\mathbb{H}}/\hat{\mathbb{I}}$, so that $(\mathbb{H} \otimes \mathbb{G}) \otimes \hat{\mathbb{G}}$ is simple and strongly coherent by Lemma 3.25);
- $\hat{\mathbb{G}}/\hat{\mathbb{I}}$ is fully symmetric over $\hat{\mathbb{I}}$ (cf. Definition 3.14): all symmetries of $\hat{\mathbb{I}}$ (i.e. of the Cayley graph of \mathbb{G}) extend to symmetries of the groupoid $\hat{\mathbb{G}}$;
- $\tilde{\mathbb{H}}/\tilde{\mathbb{I}} = ((\mathbb{H} \otimes \mathbb{G}) \otimes \hat{\mathbb{G}})/\tilde{\mathbb{I}}$ over $\tilde{\mathbb{I}} = (\hat{G}, \hat{E})$.

Proposition 3.27. *In this situation, the Cayley graph of $\hat{\mathbb{G}}$ also carries the structure of an amalgamation pattern over the coarser incidence pattern \mathbb{I} , which is complete and induces a groupoid $\tilde{\mathbb{G}}/\mathbb{I}$ (by groupoidal action on $\hat{\mathbb{G}}$ over \mathbb{I} , in the sense of Example 3.8) such that*

- (a) $\tilde{\mathbb{G}}/\mathbb{I}$ is simple and compatible with \mathbb{H}/\mathbb{I} ;
- (b) $\tilde{\mathbb{G}}/\mathbb{I}$ is fully symmetric over \mathbb{H}/\mathbb{I} if \mathbb{G} is fully symmetric over \mathbb{H}/\mathbb{I} ;
- (c) $\tilde{\mathbb{G}}/\mathbb{I}$ is N -acyclic if $\hat{\mathbb{G}}/\hat{\mathbb{I}}$ is;
- (d) the direct product $\mathbb{H} \otimes \tilde{\mathbb{G}}$ of \mathbb{H}/\mathbb{I} with this $\tilde{\mathbb{G}}/\mathbb{I}$ is a covering of \mathbb{H}/\mathbb{I} by a simple and strongly coherent amalgamation pattern;
- (e) the covering by $\mathbb{H} \otimes \tilde{\mathbb{G}}$ is fully symmetric over \mathbb{H}/\mathbb{I} if \mathbb{G} is fully symmetric over \mathbb{H}/\mathbb{I} .

In the rôle of a simple and strongly coherent covering of \mathbb{H}/\mathbb{I} , the direct product $\mathbb{H} \otimes \tilde{\mathbb{G}}$ of \mathbb{H}/\mathbb{I} with this $\tilde{\mathbb{G}}/\mathbb{I}$ can thus replace the nested direct product $(\mathbb{H} \otimes \mathbb{G}) \otimes \hat{\mathbb{G}}$ considered above. By combination with Lemmas 2.20, a (fully symmetric) realisation of \mathbb{H} can therefore also be obtained as a reduced product of \mathbb{H} with a suitable finite groupoid in a single step.

Corollary 3.28. *For any $N \geq 2$, any finite amalgamation pattern \mathbb{H}/\mathbb{I} possesses a covering by a finite, simple and strongly coherent amalgamation pattern $\tilde{\mathbb{H}}/\tilde{\mathbb{I}}$ that is fully symmetric over \mathbb{H}/\mathbb{I} , and which is obtained as a direct product of \mathbb{H}/\mathbb{I} with a suitable N -acyclic finite groupoid $\tilde{\mathbb{G}}/\mathbb{I}$.*

Proof of Proposition 3.27. We start with the construction of the groupoid $\tilde{\mathbb{G}}/\mathbb{I}$, which is obtained by re-interpreting the Cayley graph of $\hat{\mathbb{G}}$ as an amalgamation pattern over the coarser incidence pattern \mathbb{I} rather than the built-in $\hat{\mathbb{I}}$. Recall that $\hat{\mathbb{I}}$ is based on the Cayley graph of \mathbb{G} , and related to \mathbb{I} by an unbranched covering $\pi: \hat{\mathbb{I}} \rightarrow \mathbb{I}$ that maps g to $\iota_2(g)$.

As \mathbb{G} is simple we may identify $e^{\mathbb{G}}$ with e and think of E as a subset $E \subseteq G$. To cast the Cayley graph of $\hat{\mathbb{G}}$ as an induced amalgamation pattern over \mathbb{I} , we merge pieces of the finer partition according to $\hat{\mathbb{I}}$ to fit the coarser pattern \mathbb{I} . This coarsening is induced by π , which maps $g \in G$ (as a sort label of $\hat{\mathbb{I}}$) to $\iota_2(g) \in S$, $e[g] = (g, ge)$ (as a link label of $\hat{\mathbb{I}}$) to $e \in E$, and correspondingly associates $\hat{g} \in \hat{G}[* , g]$ with sort $s = \iota_2(g) \in S$. The resulting amalgamation pattern (presented both in proper multi-sorted format and as a labelled tuple) is

$$\hat{\mathbb{G}}/\mathbb{I} := (\mathbb{I}, \hat{G}, \delta: \hat{G} \rightarrow S, P, \eta: P \rightarrow E) = (\hat{G}, (\hat{G}_s)_{s \in S}, (\hat{\rho}_e)_{e \in E})$$

where δ partitions \hat{G} into the sets

$$\hat{G}_s = \bigcup \{ \hat{G}[* , g] : \iota_2(g) = s \}$$

and each $\hat{\rho}_e \in P$ is the disjoint union of the local bijections induced by right multiplications with generators $\hat{e} = e[g] = (g, ge) \in \hat{E}$, for which $\pi(\hat{e}) = e$. Since $\pi: \hat{\mathbb{I}} \rightarrow \mathbb{I}$ is an unbranched covering, there always is a unique such \hat{e} applicable at every $\hat{g} \in \hat{G}_s$, dependent on just the partition set $\hat{G}_g = \hat{G}[* , g]$ to which \hat{g} belongs in $\hat{\mathbb{G}}/\hat{\mathbb{I}}$. More specifically, the lift of $e \in E$ with $\iota(e) = (s, s')$ in \mathbb{I} to some $\hat{g} \in \hat{g} \in \hat{G}[* , g] \subseteq \hat{G}_s$ with $\iota_2(g) = s$ is $e[g] = (g, ge)$. In terms of its operation by right multiplication on the Cayley graph of $\hat{\mathbb{G}}$,

$$\hat{\rho}_e := \bigcup_{\pi(\hat{e})=e} \rho_{\hat{e}}^{\hat{\mathbb{G}}} = \bigcup_{g \in G[* , \iota_1(e)]} (\rho_{e[g]}^{\hat{\mathbb{G}}} : \hat{G}_{\iota_1(e)} \longrightarrow \hat{G}_{\iota_2(e)}),$$

each $\hat{\rho}_e$ is a bijection between the partition sets ($\hat{\mathbb{G}}/\mathbb{I}$ is complete). Extending this analysis to walks $w = e_1 \cdots e_n$ from $\iota_1(w)$ to $\iota_2(w)$ in \mathbb{I} , we write $w[g]$ for the unique lifting of the generator sequence $e_1 \cdots e_n$ that starts with $e_1[g]$ (at any $\hat{g} \in \hat{G}[* , g]$ for $g \in G[* , \iota_1(e)]$) and, writing $w_i := e_1 \cdots e_i$ for the prefixes, produces the walk

$$w[g] = e_1[g] \cdot e_2[ge_1^{\mathbb{G}}] \cdots e_n[ge_{n-1}^{\mathbb{G}}] \in \hat{\mathbb{I}}^*.$$

We obtain $\hat{\rho}_w$ as the composition of the $\hat{\rho}_{e_i}$ along this walk,

$$\hat{\rho}_w = \bigcup_{g \in G[* , \iota_1(w)]} (\rho_{w[g]}^{\hat{\mathbb{G}}} : \hat{G}_{\iota_1(w)} \longrightarrow \hat{G}_{\iota_2(w)}),$$

which is again a bijection between the partition sets involved, and maps every $\hat{g} \in \hat{G}[* , g] \subseteq \hat{G}_s$ (for $g \in G[* , s]$, $s = \iota_1(e_1)$) according to

$$\hat{\rho}_w : \hat{g} \longmapsto \hat{g} \cdot e_1[g] \cdot e_2[ge_1^{\mathbb{G}}] \cdots e_n[ge_{n-1}^{\mathbb{G}}].$$

We let $\tilde{\mathbb{G}}/\mathbb{I}$ be the groupoid generated by this action of the $\rho_e^{\hat{\mathbb{G}}}$ for $e \in E$, in the inverse semigroup $I(\hat{G})$. Note how this groupoid structure is induced in terms of the derived groupoidal action of \mathbb{I}^* on the Cayley graph of $\hat{\mathbb{G}}$, which combines in parallel the effect of all the $\rho_{e[g]}$ for $g \in G[* , \iota_1(e)]$ according to

the native groupoidal action of $\hat{\mathbb{I}}^*$ on $\hat{\mathbb{G}}$. So the elements of $\tilde{\mathbb{G}}$ are the partial bijections $\hat{\rho}_w$ of \hat{G} induced by walks w in \mathbb{I} .

For a walk w from s to s in \mathbb{I} , $\hat{\rho}_w$ represents 1_s in the groupoid $\tilde{\mathbb{G}}$ induced by this action, if, and only if, $\hat{\rho}_w$ is the identity of \hat{G}_s , i.e. if, and only if, $\rho_{w[g]}^{\hat{G}} = \text{id}_{\hat{G}[* , g]}$ for all $g \in G[* , s]$, if, and only if, right multiplication with $(w[g])^{\hat{G}}$ is a local identity on $\hat{G}[* , g]$ for all $g \in G[* , s]$. But any two $g, g' \in G[* , s]$ are related by a unique \mathbb{I} -rigid automorphism σ of $\hat{\mathbb{I}}$ (i.e. the Cayley graph of \mathbb{G} , not the groupoid) that maps g to g' . Due to the assumption of full symmetry of $\tilde{\mathbb{G}}$ over $\hat{\mathbb{I}}$, this automorphism extends to a symmetry $\hat{\sigma}$ of the groupoid $\tilde{\mathbb{G}}$, and this extension automorphically relates the actions of $\rho_{w[g]}^{\hat{G}}$ at $\hat{G}[* , g]$ and $\rho_{w[g']}^{\hat{G}}$ at $\hat{G}[* , g']$, as indicated in the following commuting diagram:

$$\begin{array}{ccccccccccc}
G[* , s] & & g & & e & & e[g] & & w[g] & & \hat{G}[* , g] & \xrightarrow[\rho_{w[g]}^{\hat{G}}]{\hat{\rho}_w} & \hat{G}[* , gw^{\mathbb{G}}] \\
\downarrow \sigma & & \downarrow \sigma & & \downarrow \sigma & & \downarrow \hat{\sigma} & & \downarrow \hat{\sigma} & & \downarrow \hat{\sigma} & & \downarrow \hat{\sigma} \\
G[* , s] & & g' & & e & & e[g'] & & w[g'] & & \hat{G}[* , g'] & \xrightarrow[\rho_{w[g']}^{\hat{G}}]{\hat{\rho}_w} & \hat{G}[* , g'w^{\mathbb{G}}]
\end{array}$$

So the following are equivalent:

- (i) the walk $w = e_1 \cdots e_n$ in \mathbb{I} from s to s generates the unit 1_s in the groupoid $\tilde{\mathbb{G}}/\mathbb{I}$: $w^{\tilde{\mathbb{G}}} = 1_s$.
- (ii) the lift $w[g]$ of this walk to a walk at any $\hat{g} \in \hat{G}[* , g] \subseteq \hat{G}_s$ in the Cayley graph of $\tilde{\mathbb{G}}$ loops back to \hat{g} , i.e. $(w[g])^{\tilde{\mathbb{G}}}$ acts as the identity of $\hat{G}[* , g]$, for all $g \in G[* , s]$.
- (iii) the lift $w[g]$ of this walk to a walk at \hat{g} loops back to \hat{g} for some $\hat{g} \in \hat{G}[* , g] \subseteq \hat{G}_s$ and some $g \in G[* , s]$.

For claim (a) of the proposition it is immediate that $\tilde{\mathbb{G}}/\mathbb{I}$ is simple, by simplicity of \mathbb{G} ; compatibility with \mathbb{H}/\mathbb{I} follows from compatibility of \mathbb{G} with \mathbb{H}/\mathbb{I} : any $\hat{\rho}_w$ that represents a unit 1_s in $\tilde{\mathbb{G}}$ must, by construction of $\tilde{\mathbb{G}}$, act as the identity of \hat{G}_s in the amalgamation pattern $\tilde{\mathbb{G}}/\mathbb{I}$, so that in particular $w^{\tilde{\mathbb{G}}} = 1_s$ in $\tilde{\mathbb{G}}$.

For (b), full symmetry of $\tilde{\mathbb{G}}/\mathbb{I}$ over \mathbb{H}/\mathbb{I} is immediate if \mathbb{G} itself is fully symmetric over \mathbb{H}/\mathbb{I} : any automorphism of \mathbb{H}/\mathbb{I} lifts to a groupoid automorphism of \mathbb{G} , which in particular induces an automorphism of its Cayley graph, and hence of the incidence pattern $\hat{\mathbb{I}}$, which in turn lifts to an automorphism of $\tilde{\mathbb{G}}$ and thus of the groupoid $\tilde{\mathbb{G}}/\mathbb{I}$ (whose definition in terms of $\tilde{\mathbb{G}}$ refers just to the Cayley graph of $\tilde{\mathbb{G}}$ and partitions relative to \mathbb{I}).

For (c), we argue that any coset cycle in $\tilde{\mathbb{G}}/\mathbb{I}$ induces a coset cycle of the same length in $\tilde{\mathbb{G}}/\hat{\mathbb{I}}$. The relevant cosets of $\tilde{\mathbb{G}}$ are generated by sub-groupoids of the form $\tilde{\mathbb{G}}[\alpha]$ for generator sets $\alpha = \alpha^{-1} \subseteq E$; with this sub-groupoid of $\tilde{\mathbb{G}}$ we

can associate the sub-groupoid $\hat{\mathbb{G}}[\hat{\alpha}]$ of $\hat{\mathbb{G}}$, generated by $\hat{\alpha} := \pi^{-1}(\alpha) = \{e[g] = (g, ge) \in \hat{E} : e \in \alpha, g \in G[* , \iota_1(e)]\}$. It remains to check that a cyclic tuple of pointed cosets $(\tilde{g}_i, \tilde{g}_i \tilde{\mathbb{G}}[\alpha_i])_{i \in \mathbb{Z}_n}$ that forms a coset cycle in $\tilde{\mathbb{G}}$, for suitable choices of the \tilde{g}_i in $\tilde{\mathbb{G}}$ translates into a cyclic tuple $(\hat{g}_i, \hat{g}_i \hat{\mathbb{G}}[\hat{\alpha}_i])_{i \in \mathbb{Z}_n}$, that forms a coset cycle in $\hat{\mathbb{G}}/\hat{\mathbb{I}}$ (cf. conditions (i) and (ii) from Definition 3.17):

(i) From $\tilde{g}_{i+1} \in g_i \tilde{\mathbb{G}}[\alpha_i]$ we find $w_i \in \alpha_i^*$ such that $\tilde{g}_{i+1} = \tilde{g}_i w_i^{\tilde{\mathbb{G}}}$ and put $\hat{g}_{i+1} := \hat{g}_i \cdot (w_i[g_i])^{\hat{\mathbb{G}}}$. Here $w_i[g_i]$ is the unique lift of the walk w_i in \mathbb{I} to a walk in $\hat{\mathbb{I}}$ at g_i if $\hat{g}_i \in \hat{G}[* , g_i]$, and consists of links in $\hat{\alpha}_i$. The choice of a starting element, say \hat{g}_0 , is arbitrary, but it is important that tracing the sequence of the $w_i[g_i]$ in $\tilde{\mathbb{G}}$ takes us back to that starting point — which it does because the product/concatenation of the w_i represents a unit in $\tilde{\mathbb{G}}$, which implies the same for their lifts in $\hat{\mathbb{G}}$.

(ii) From $\tilde{g}_i \tilde{\mathbb{G}}[\alpha_i \cap \alpha_{i-1}] \cap \tilde{g}_{i+1} \tilde{\mathbb{G}}[\alpha_i \cap \alpha_{i+1}] = \emptyset$, we directly infer the corresponding condition for the translation to $\hat{\mathbb{G}}$ as follows. If \hat{g} were in the corresponding intersection in $\hat{\mathbb{G}}$, it could be represented as $\hat{g} = \hat{g}_i \hat{u}^{\hat{\mathbb{G}}} = \hat{g}_{i+1} \hat{w}^{\hat{\mathbb{G}}}$ for walks $\hat{u} \in (\hat{\alpha}_i \cap \hat{\alpha}_{i-1})^*$ and $\hat{w} \in (\hat{\alpha}_i \cap \hat{\alpha}_{i+1})^*$ in $\hat{\mathbb{I}}$. These walks project to walks $u = \pi(\hat{u})$ and $w = \pi(\hat{w})$ in \mathbb{I} , which would show that also $\tilde{g}_i \tilde{\mathbb{G}}[\alpha_i \cap \alpha_{i-1}] \cap \tilde{g}_{i+1} \tilde{\mathbb{G}}[\alpha_i \cap \alpha_{i+1}] \neq \emptyset$.

For (d), the direct product $\tilde{\mathbb{H}}' := \mathbb{H} \otimes \tilde{\mathbb{G}}$ is a covering of \mathbb{H}/\mathbb{I} by definition; it remains to show that $\tilde{\mathbb{H}}'/\tilde{\mathbb{I}}'$ is simple and strongly coherent as an amalgamation pattern over the incidence pattern $\tilde{\mathbb{I}}'$ induced by the Cayley graph of $\tilde{\mathbb{G}}$. Recall from Definition 2.8 that simplicity requires the partial isomorphisms $\rho_{\tilde{e}}$ in $\tilde{\mathbb{H}}'$ to be maximal among all ρ_w in $\tilde{\mathbb{H}}'$ along walks \tilde{w} from $\iota_1(\tilde{e})$ to $\iota_2(\tilde{e})$ in $\tilde{\mathbb{I}}'$. The links \tilde{e} of $\tilde{\mathbb{I}}'$ correspond to right multiplication in $\tilde{\mathbb{G}}$ by generators $e^{\tilde{\mathbb{G}}}$ at $\tilde{g} \in \tilde{G}$, which in $\tilde{\mathbb{H}}' = \mathbb{H} \otimes \tilde{\mathbb{G}}$ act as

$$\rho_{\tilde{e}} : (h, \tilde{g}) \longmapsto (\rho_e(h), \tilde{g} e^{\tilde{\mathbb{G}}})$$

on $h \in A_s$ for $s = \iota_1(e)$. Any walk \tilde{w} from $\tilde{g} = \iota_1(\tilde{e})$ to $\tilde{g}' = \iota_2(\tilde{e})$ in $\tilde{\mathbb{I}}'$ is such that $\tilde{w}^{\tilde{\mathbb{G}}} = e^{\tilde{\mathbb{G}}}$ in $\tilde{\mathbb{G}}$. Since $\tilde{\mathbb{G}}$ is a groupoid over \mathbb{I} , this implies in particular that, for the projection w of \tilde{w} to \mathbb{I} , $w^{\mathbb{G}} = e$ in \mathbb{G} , whence $\rho_w(h) = \rho_e(h)$ (where defined) is guaranteed by compatibility of \mathbb{G} with \mathbb{H} .

Strong coherence, by Definition 2.8, requires that, for any two walks \tilde{w}_1, \tilde{w}_2 from \tilde{g} to \tilde{g}' in $\tilde{\mathbb{I}}'$, there is a walk \tilde{w} such that $\rho_{\tilde{w}}$ in $\tilde{\mathbb{H}}'$ is a common extension of the $\rho_{\tilde{w}_i}$ for $i = 1, 2$. The reason essentially is that $\tilde{w}_1^{\tilde{\mathbb{G}}} = \tilde{w}_2^{\tilde{\mathbb{G}}}$ implies that the actions of the \tilde{w}_i on $((\mathbb{H} \otimes \mathbb{G}) \otimes \tilde{\mathbb{G}})$ agree, where strong coherence yields a common extension. More specifically, consider the effect of the actions of \tilde{e} (links in $\tilde{\mathbb{I}}'$) along the walks \tilde{w}_i in $\tilde{\mathbb{I}}'$ in the groupoid component $\tilde{\mathbb{G}}$ of $\tilde{\mathbb{H}}' = \mathbb{H} \otimes \tilde{\mathbb{G}}$. The exploration of $\tilde{\mathbb{G}}$ given above shows that $\tilde{w}_1^{\tilde{\mathbb{G}}} = \tilde{w}_2^{\tilde{\mathbb{G}}}$ implies that the projections of the walks \tilde{w}_i to walks w_i in \mathbb{I} induce the same action on $\tilde{\mathbb{G}}$. In terms of matching walks $w_i[g]$ in the Cayley graph of $\tilde{\mathbb{G}}$, which are the lifts of the w_i at appropriate elements $\hat{g} \in \hat{\mathbb{G}}[* , g]$, this implies that $(w_1[g])^{\tilde{\mathbb{G}}} = (w_2[g])^{\tilde{\mathbb{G}}}$. Strong coherence of $((\mathbb{H} \otimes \mathbb{G}) \otimes \tilde{\mathbb{G}})$ therefore implies the existence of a walk $w[g]$ in the Cayley graph of $\tilde{\mathbb{G}}$ whose action on $((\mathbb{H} \otimes \mathbb{G}) \otimes \tilde{\mathbb{G}})$ extends those of the $w_i[g]$.

Now $(w[g])^{\hat{\mathbb{G}}} = (w_i[g])^{\hat{\mathbb{G}}}$ for all appropriate g implies that also $w^{\hat{\mathbb{G}}} = \tilde{w}^{\hat{\mathbb{G}}}$, since these groupoid elements are defined in term of the action on $\hat{\mathbb{G}}$. Since their action in the \mathbb{H} -component of $\hat{\mathbb{H}}' = \mathbb{H} \otimes \hat{\mathbb{G}}$ is fully determined by the projections w and w_i to \mathbb{I} in both $((\mathbb{H} \otimes \mathbb{G}) \otimes \hat{\mathbb{G}})$ and $\mathbb{H} \otimes \hat{\mathbb{G}}$, $\rho_{w_i} \subseteq \rho_w$ in \mathbb{H} implies the same also over $\mathbb{H} \otimes \hat{\mathbb{G}}$.

Claim (e) follows from (b) together with Lemma 3.24. \square

4 Generic realisations in reduced products

Recall the equivalence relation \approx defined in connection with Lemma 2.10: it is the equivalence relation induced on the universe $H = \bigcup_{s \in S} A_s$ of an amalgamation pattern $\mathbb{H} = (\mathcal{H}, (\mathcal{A}_s)_{s \in S}, (\rho_e)_{e \in E})$ over $\mathbb{I} = (S, E)$ by regarding as equivalent any elements $a \in A_s$ and $\rho_e(a) \in A_{s'}$ for links $e \in E$ with $\iota(e) = (s, s')$. Also recall the notion of an atlas from Definition 2.4. The \approx -equivalence class of $a \in H$ is denoted as $[a]$.

Definition 4.1. Let \mathbb{G}/\mathbb{I} be compatible with \mathbb{H}/\mathbb{I} . The *reduced product* $\mathbb{H} \otimes \mathbb{G}/\approx$ of an amalgamation pattern \mathbb{H}/\mathbb{I} with a groupoid \mathbb{G}/\mathbb{I} is based on the quotient structure of the relational structure $\hat{\mathcal{H}}$ of the direct product $\hat{\mathbb{H}} = \mathbb{H} \otimes \mathbb{G}$ w.r.t. the equivalence relation \approx whenever this quotient is well-defined.¹⁰ We endow this relational structure $\hat{\mathcal{H}}/\approx$ with an atlas induced by the families of subsets

$$U_s := \{u[g] : g \in G[* , s]\} \text{ where } u[g] = [A_{\iota_2(g)} \times \{g\}] = \{[(a, g)] : a \in A_{\iota_2(g)}\}$$

for $s \in S$, with the natural isomorphisms

$$\begin{array}{ccc} \pi_{u,s} : \hat{\mathcal{H}}/\approx & \rightarrow & \mathcal{A}_s \\ [(a, g)] & \mapsto & a \end{array}$$

for $u = u[g] \in U_s$, which is well-defined due to coherence of the product. With this atlas on $U = \bigcup_{s \in S} U_s$ over the relational structure $\hat{\mathcal{H}}/\approx$:

$$\mathbb{H} \otimes \mathbb{G}/\approx = (\hat{\mathcal{H}}/\approx, U, (U_s)_{s \in S}, (\pi_{u,s})_{u \in U_s}).$$

It is clear from the format of these reduced products in relation to \mathbb{H}/\mathbb{I} that they are natural candidates for realisations of \mathbb{H}/\mathbb{I} according to Definition 2.14. And indeed, whenever the direct product $\mathbb{H} \otimes \mathbb{G}$ is simple and strongly coherent, Lemma 2.20 shows that $\mathbb{H} \otimes \mathbb{G}/\approx$ is a realisation of $\mathbb{H} \otimes \mathbb{G}$ which induces a realisation of \mathbb{H}/\mathbb{I} by Remark 2.19. So Corollary 3.28 guarantees the existence of finite realisations of \mathbb{H}/\mathbb{I} based on reduced products of \mathbb{H}/\mathbb{I} with suitable finite groupoids \mathbb{G}/\mathbb{I} .

Theorem 4.2. *Every finite amalgamation pattern \mathbb{H}/\mathbb{I} admits finite realisations by natural reduced products $\mathbb{H} \otimes \mathbb{G}/\approx$ based on a direct product $\mathbb{H} \otimes \mathbb{G}$ of \mathbb{H} with a*

¹⁰I.e. whenever the direct product $\hat{\mathbb{H}}$ is globally consistent, which is in particular the case if the direct product is strongly coherent, by Lemma 2.10.

suitable finite groupoid \mathbb{G}/\mathbb{I} , such that $\mathbb{H} \otimes \mathbb{G}$ itself is a simple, strongly coherent covering of \mathbb{H} . Such realisations can be chosen to be fully symmetric over the given amalgamation pattern \mathbb{H}/\mathbb{I} , and \mathbb{G}/\mathbb{I} can be chosen to be N -acyclic for any given $N \geq 2$.

Remark 4.3. *The canonical infinite realisation of \mathbb{H}/\mathbb{I} of Remark 3.2 can also be cast as a reduced product of \mathbb{H} with the free groupoid over \mathbb{I} of Remark 3.9.*

4.1 Degrees of acyclicity in realisations

As seen above, Corollary 3.28 together with Lemma 2.20 offer a route to finite realisations with additional acyclicity properties, obtained as reduced products with groupoids that are N -acyclic for any chosen $N \geq 2$. We recall that 2-acyclicity is essential towards strong coherence. We now want to analyse the effect that higher levels of groupoid coset acyclicity may have on the resulting realisation. To this end we first review an established notion of hypergraph acyclicity ([6], also known as α -acyclicity in the literature, e.g. [5]). Graded, i.e. quantitatively localised or size bounded versions of this natural notion can be applied to the hypergraph structure of the atlas of a finite realisation. Realisations obtained as reduced products with suitable N -acyclic groupoids can then be shown to have an atlas that is N -acyclic in the sense of hypergraph acyclicity. This means that the family of distinguished substructures induced by the coordinate domains, or the isomorphic copies of the template structures \mathcal{A}_s that make up the realisation locally overlap in a tree-like (viz. tree-decomposable) fashion. This will further imply graded local universality properties for the realisation.

4.2 Hypergraph acyclicity

Definition 4.4. A *hypergraph* (A, U) is a structure consisting of a *vertex set* A and a *set of hyperedges* $U \subseteq \mathcal{P}_{\text{fin}}(A)$ such that $A = \bigcup U$.¹¹

With a hypergraph (A, U) , with set of hyperedges $U \subseteq \mathcal{P}_{\text{fin}}(A)$ over vertex set A , we associate its *Gaifman graph*, which is a simple undirected graph over the same vertex set. Its edge relation R links two distinct vertices a and a' if they occur in a common hyperedge:

$$R = \{(a, a') \in A^2 : a \neq a', a, a' \in u \text{ for some } u \in U\}.$$

In other words, the Gaifman graph consists of the union of cliques over the $u \in U$.

Recall that a *chord* in a cycle $(a_i)_{i \in \mathbb{Z}_n}$ is an edge linking two vertices that are not next neighbours in the cycle (which can only occur in cycles of length greater than 3).

¹¹Hyperedges are always required to be finite subsets of the set of vertices even where this set is infinite; forbidding vertices that are outside all hyperedges is essentially w.l.o.g. (noting that singleton hyperedges are allowed) but helps to avoid trivial special provisions in some constructions and arguments.

Definition 4.5. A hypergraph (A, U) is *acyclic* if it is both *chordal* and *conformal*, i.e. if

- (i) its Gaifman graph has no chordless cycles of length greater than 3 (chordality);
- (ii) every clique in its Gaifman graph is fully contained in some hyperedge (conformality).

Useful equivalent characterisations (cf. [5]) in the case of finite hypergraphs involve the notion of *tree-decomposability* and a notion of *retractability*.

Definition 4.6. A *tree decomposition* $(T, E, \lambda: T \rightarrow U)$ of a finite hypergraph (A, U) consists of a labelling of the hyperedges $u \in U$ by the vertices of a finite (graph-theoretic) tree (T, E) via a surjective function $\lambda: T \rightarrow U$ such that, for every $a \in A$, the subset $\{t \in T: a \in \lambda(t)\} \subseteq T$ is connected in (T, E) .

A finite hypergraph (A, U) is *retractable* (in the sense of Graham's algorithm) if it can be reduced to the empty structure by recursive application of the following retraction steps:

- deletion from A and every $u \in U$ of a single vertex $a \in A$ that is an element of just one hyperedge;
- deletion from U of a single hyperedge that is a subset of some other hyperedge.

It is obvious that a hypergraph is acyclic if, and only if, all its finite induced sub-hypergraphs are acyclic. For finite hypergraphs it is also well known, e.g. from [5], that acyclicity coincides with the existence of a tree decomposition as well as with retractability. Rather than outright acyclicity, we shall focus on a notion of N -acyclicity as a suitable quantitative approximation to acyclicity, so that acyclicity becomes the limit across all levels of N -acyclicity.

Definition 4.7. For $N \geq 3$, a hypergraph (A, U) is *N -acyclic* if every induced sub-hypergraph with up to N many vertices is acyclic.

It is easy to check that N -acyclicity is equivalent to the combination of

- (i) N -chordality: its Gaifman graph has no chordless cycles of lengths greater than 3 and up to N ;
- (ii) N -conformality: every clique of size up to N in its Gaifman graph is contained in some hyperedge.

In the following we also refer to *N -acyclic atlases* or to *N -acyclic realisations*, meaning that the hypergraph associated with the co-ordinate domains of the atlas is N -acyclic.

Observation 4.8. *If (A, U) is the hypergraph of co-ordinate domains of an atlas of a finite relational structure \mathcal{A} on A according to Definition 2.4, then a tree decomposition $(T, E, \lambda: T \rightarrow U)$ of (A, U) is such that every tuple in a relation of \mathcal{A} must be fully contained in one of the $\lambda(t)$ and such that all elements in the overlap $\lambda(t_1) \cap \lambda(t_2)$ of two charts must be represented in every $\lambda(t)$ along the shortest connecting path between t_1 and t_2 in the tree T . Model-theoretically,*

such a tree decomposition provides a representation of the underlying structure as a free amalgam of its distinguished substructures $\mathcal{A} \upharpoonright u$ for $u \in U$. The free amalgam can be put together in a stepwise fashion by following in reverse order a retraction as in Definition 4.6.

4.3 Acyclicity in reduced products

Compare Definition 4.1 for reduced products and recall that compatibility of the groupoid \mathbb{G}/\mathbb{I} with the amalgamation pattern \mathbb{H}/\mathbb{I} is equivalent with the coherence of the direct product $\mathbb{H} \otimes \mathbb{G}$. The strictly stronger assumption of strong coherence of the direct product $\mathbb{H} \otimes \mathbb{G}$ implies that the reduced product $\mathbb{H} \otimes \mathbb{G}/\approx$ is well-defined. The next lemma refers to the hypergraph induced by the natural atlas of this reduced product $\mathbb{H} \otimes \mathbb{G}/\approx$, whose hyperedges are the co-ordinate domains of its charts onto the \mathcal{A}_s of \mathbb{H}/\mathbb{I} ,

$$u[g] := [A_s \times \{g\}] = \{[(a, g)]: a \in A_s\},$$

for all $g \in G[*], s \in S$ (note that $g \in \mathbb{G}$ fully determines $u[g]$ since $s = \iota_2(g)$). Note that an element $[(a, g)]$ of the reduced product, which is represented by a pair (a, g) of the direct product with $a \in A_s$ and $g \in G[*], s$ for some $s \in S$ (which thus is an element of $u[g]$), is also an element of $u[h]$ for some $h \in \mathbb{G}[*], s'$ if, and only if, it is represented by some pair $(a', h) \approx (a, g)$ for some $a' \in A_{s'}$. This is the case if, and only if, there is a walk w from s to s' in \mathbb{I} such that $h = gw^{\mathbb{G}}$ in \mathbb{G} and $\rho_w(a) = a'$ in \mathbb{H} . We observe that any such w must be comprised of generators $e \in E$ for which $\rho_{w'}(a) \in \text{dom}(\rho_e)$ for some w' from s to $\iota_1(e)$ for which $a \in \text{dom}(\rho_{w'})$, so that $(a, g) \approx (\rho_{w'}(a), gw'^{\mathbb{G}}) \approx (\rho_{w'e}(a), g(w'e)^{\mathbb{G}})$. It will therefore be convenient to define these generator sets as

$$\alpha_a = \{e \in E: \rho_w(a) \in \text{dom}(\rho_e) \text{ for some walk } w \text{ from } s \text{ to } \iota_1(e) \text{ in } \mathbb{I}\}.$$

Note that $\alpha_a^{-1} = \alpha_a$ is closed under edge reversal (inversion of generators) since $\rho_w(a) \in \text{dom}(\rho_e)$ if, and only if, $\rho_{w_e}(a) \in \text{dom}(\rho_e^{-1}) = \text{dom}(\rho_{e^{-1}})$. We can isolate the important connection between cosets w.r.t. these generator sets and the hypergraph structure of the atlas of $\mathbb{H} \otimes \mathbb{G}/\approx$ as follows: $g\mathbb{G}[\alpha_a]$ is the set of groupoid elements at which $[(a, g)]$ is represented.

Observation 4.9. *In any reduced product $\mathbb{H} \otimes \mathbb{G}/\approx$ (which by definition is based on a coherent direct product), for $s \in S$, any $a \in A_s \subseteq H$ and $g \in G[*], s$:*

$$\{h \in \mathbb{G}: [(a, g)] \in u[h]\} = g\mathbb{G}[\alpha_a].$$

For N -acyclicity of \mathbb{G}/\mathbb{I} see Definition 3.18 in connection with the following.

Lemma 4.10. *For any amalgamation pattern \mathbb{H}/\mathbb{I} and finite groupoid \mathbb{G}/\mathbb{I} such that the direct product $\mathbb{H} \otimes \mathbb{G}$ is strongly coherent: if \mathbb{G}/\mathbb{I} is N -acyclic for some $N \geq 3$, then the hypergraph induced by the co-ordinate domains $u[g]$ of the natural atlas the reduced product $\mathbb{H} \otimes \mathbb{G}/\approx$ is N -acyclic.*

Proof. Denote as (A, U) the hypergraph formed by the universe A of the relational structure \mathcal{A} of the reduced product $\mathbb{H} \otimes \mathbb{G}/\approx$ together with the set of co-ordinate domains of its atlas. Its vertex set A consists of the \approx -equivalence classes $[(a, g)]$ of elements (a, g) of the direct product and its hyperedges are the sets

$$u[g] := [A_{\iota_2(g)} \times \{g\}] = \{[(a, g)]: a \in A_{\iota_2(g)}\}.$$

Assume that \mathbb{G}/\mathbb{I} is N -acyclic. We show that (A, U) is N -acyclic.

N-chordality. For N -chordality, we need to show that any chordless cycle in the Gaifman graph of (A, U) must have length greater than N . It suffices to argue that any chordless cycle in (A, U) induces a coset cycle of the same length in \mathbb{G} . Let $(c_i)_{i \in \mathbb{Z}_n}$ be a chordless cycle in the Gaifman graph of (A, U) : for each $i \in \mathbb{Z}_n$ there is some hyperedge $u_i \in U$ such that $c_{i-1}, c_i \in u_i$, while there is no such $u \in U$ for the pair c_i, c_j if $j \neq i \pm 1$ in \mathbb{Z}_n . Fix a sequence of such $u_i = u[g_i]$ so that $c_i \in u_i, u_{i+1}$. Since $c_i \in u_i$ we may choose representatives such that $c_i = [(a_i, g_i)]$. We associate with each $[(a_i, g_i)]$ the generator set $\alpha_i = \alpha_{a_i} \subseteq E$ discussed in Observation 4.9:

$$\alpha_i = \{e \in E: \rho_w(a_i) \in \text{dom}(\rho_e) \text{ for some walk } w \text{ from } s_i \text{ to } \iota_1(e) \text{ in } \mathbb{I}\}.$$

With $i \in \mathbb{Z}_n$ we associate the sub-groupoid $\mathbb{G}[\alpha_i]$ and the pointed coset $(g_i, g_i\mathbb{G}[\alpha_i])$, and show that

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

forms a coset cycle in \mathbb{G}/\mathbb{I} by checking the two defining conditions of Definition 3.17:

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

For (i), observe that $c_i = [(a_i, g_i)] \in u_{i+1} = u[g_{i+1}]$ implies that there is some walk w from s_i to s_{i+1} in \mathbb{I} such that $g_{i+1} = g_i w^{\mathbb{G}}$ in \mathbb{G} and such that $a_i \in A_{s_i}$ is mapped to some element of $A_{s_{i+1}}$ by ρ_w in \mathbb{H}/\mathbb{I} (cf. Observation 4.9). It follows that this path is made up of links $e \in \alpha_i$ so that $w^{\mathbb{G}} \in \mathbb{G}[\alpha_i]$ whence $g_{i+1} \in g_i\mathbb{G}[\alpha_i]$.

The argument for (ii) relies on the fact that the cycle formed by the c_i in (A, U) is chordless: suppose that, to the contrary of (ii), there were some groupoid element $g \in g_i\mathbb{G}[\alpha_i \cap \alpha_{i-1}] \cap g_{i+1}\mathbb{G}[\alpha_i \cap \alpha_{i+1}]$. By Observation 4.9, $c_j = [(a_j, g_j)] \in u[g]$ if $g \in g_j\mathbb{G}[\alpha_j]$. So obviously $c_i, c_{i+1} \in u[g]$. But also $c_{i-1} \in u_i = u[g_i]$ by our initial choice of the u_i , whence, by Observation 4.9, $g_i \in g_{i-1}\mathbb{G}[\alpha_{i-1}]$. Since $g \in g_i\mathbb{G}[\alpha_{i-1}]$, this implies that $g \in g_{i-1}\mathbb{G}[\alpha_{i-1}]$ too, whence, again by Observation 4.9, $c_{i-1} \in u[g]$ as well. So $u[g]$ would be a chord.

N-conformality. Towards an indirect proof of N -conformality consider a minimal clique $C \subseteq A$ in the Gaifman graph of (A, U) that is not contained in one of the hyperedges $u \in U$; minimality here means that every proper subset of C is contained in some hyperedge. It suffices to show that any such clique induces a coset cycle of length $|C|$ in \mathbb{G} . Enumerate C as $C = \{c_i: i \in \mathbb{Z}_n\}$ for $n = |C| \geq 3$

(as $|C| \geq 3$ follows from the assumptions). For $i \in \mathbb{Z}_n$ let u_i be a hyperedge $u_i = u[k_i]$ that contains $C \setminus \{c_i\}$, so that $c_i \in u_j$ if, and only if, $j \neq i$. Also fix a choice of representatives for the vertices $c_i \in C$ in the direct product $\mathbb{H} \otimes \mathbb{G}$, according to $c_i = [(a_i, g_i)]$ with $a_i \in A_{s_i}$ and $g_i \in G[* , s_i]$, and let as above

$$\alpha_i = \{e \in E : \rho_w(a_i) \in \text{dom}(\rho_e) \text{ for some walk } w \text{ from } s_i \text{ to } \iota_1(e) \text{ in } \mathbb{I}\}.$$

Further let, for $i \in \mathbb{Z}_n$, $\beta_i := \bigcap_{j \neq i, i+1} \alpha_j$ and $h_i := k_i^{-1}k_{i+1}$. We claim that

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

forms a coset cycle in \mathbb{G} , and need to establish conditions (i) and (ii) as above.

Towards (i) we note that $c_j \in u_i = u[k_i]$ for $i \neq j$, by Observation 4.9 implies that $k_i \in g_j \mathbb{G}[\alpha_j]$ for all $i \neq j$ so that in particular $k_i^{-1}k_{i+1} \in \mathbb{G}[\alpha_j]$ for all $j \neq i, i+1$. Since \mathbb{G} is 2-acyclic, Observation 3.19 implies that $k_i^{-1}k_{i+1} \in \bigcap_{j \neq i, i+1} \mathbb{G}[\alpha_j] = \mathbb{G}[\beta_i]$, or that $k_{i+1} \in k_i \mathbb{G}[\beta_i]$ as desired.

For (ii), assume to the contrary that there were some $k \in k_i \mathbb{G}[\beta_i \cap \beta_{i-1}] \cap k_{i+1} \mathbb{G}[\beta_i \cap \beta_{i+1}]$. Note that $\beta_i \cap \beta_{i-1} = \bigcap_{j \neq i} \alpha_j$ and $\beta_i \cap \beta_{i+1} = \bigcap_{j \neq i+1} \alpha_j$. We claim that $u[k]$ would contain all of C , contradicting assumptions.

Indeed, for $j \neq i$, $c_j = [(a_j, g_j)] \in u_i = u[k_i]$ so that, by Observation 4.9, $k_i \in g_j \mathbb{G}[\alpha_j]$. Since also $k \in k_i \mathbb{G}[\beta_i \cap \beta_{i-1}] \subseteq \mathbb{G}[\alpha_j]$ for all $j \neq i$, $k \in g_j \mathbb{G}[\alpha_j]$ and $c_j \in u[k]$ for $j \neq i$ follows, again by Observation 4.9. A strictly analogous argument, based on $k \in k_{i+1} \mathbb{G}[\beta_i \cap \beta_{i+1}]$ and $c_j \in u_{i+1} = u[k_{i+1}]$ for $j \neq i+1$, shows that $c_j \in u[k]$ for all $j \neq i+1$. Overall, therefore $c_j \in u[k]$ for all i , the desired contradiction. \square

4.4 Acyclicity in realisations

Together with Theorem 4.2 we therefore find that not just fully symmetric realisations but fully symmetric realisations of any degree of hypergraph acyclicity can be obtained as reduced products with suitable finite groupoids. Section 4.5 below will show why N -acyclic realisations are particularly desirable.

Corollary 4.11. *For any $N \geq 2$, every finite amalgamation pattern \mathbb{H}/\mathbb{I} admits finite realisations by natural reduced products $\mathbb{H} \otimes \mathbb{G} / \approx$ based on a direct product $\mathbb{H} \otimes \mathbb{G}$ of \mathbb{H} with a suitable finite N -acyclic groupoid \mathbb{G}/\mathbb{I} , such that $\mathbb{H} \otimes \mathbb{G}$ itself is a simple, strongly coherent covering of \mathbb{H} . Such realisations can be chosen to be fully symmetric over the given amalgamation pattern \mathbb{H}/\mathbb{I} and N -acyclic (in the sense that the hypergraph of co-ordinate domains, or of the distinguished substructures isomorphic to the \mathcal{A}_s of \mathbb{H} is N -acyclic).*

Remark 4.12. *The canonical infinite realisation of \mathbb{H}/\mathbb{I} as discussed in Remarks 3.2 and 4.3 above is N -acyclic for all N , i.e. its hypergraph of co-ordinate domains is acyclic. This can be seen directly or via its nature as a reduced product with the acyclic free groupoid over \mathbb{I} of Remarks 3.9 and 3.20.*

4.5 A universality property for N -acyclic realisations

Definition 4.13. A realisation \mathbb{A}/\mathbb{H} of an amalgamation pattern \mathbb{H}/\mathbb{I} is N -universal if every induced substructure of size up to N of \mathcal{A} can be homomorphically mapped into any other (finite or infinite) realisation of \mathbb{H}/\mathbb{I} .

Proposition 4.14. N -acyclic realisations are N -universal.

Note that the canonical infinite realisation of \mathbb{H}/\mathbb{I} (cf. Remarks 3.2 and 4.3), being N -acyclic for all N by Remark 4.12, must be N -universal for all N . In fact, it does admit full homomorphisms into any other realisation, by a countable chain argument analogous to the following finite chain argument.

Proof. For given \mathbb{H}/\mathbb{I} , let $\mathbb{A}/\mathbb{H} = (\mathbb{H}/\mathbb{I}, \mathcal{A}, U, (U_s)_{s \in S}, (\pi_{u,s})_{u \in U_s})$ be an N -acyclic realisation, $\mathbb{B}/\mathbb{H} = (\mathbb{H}/\mathbb{I}, \mathcal{B}, V, (V_s)_{s \in S}, (\pi_{v,s})_{v \in V_s})$ any other realisation. By the characterisation of hypergraph acyclicity in terms of tree decompositions from [5], we know that any induced substructure $\mathcal{C} = \mathcal{A} \upharpoonright C \subseteq \mathcal{A}$ of size $|C| \leq N$ admits a tree decomposition $(T, E, \lambda: T \rightarrow U)$ according to Definition 4.6 such that $C \subseteq \bigcup_{t \in T} \lambda(t)$ and for every $c \in C$ the set $\{t \in T: c \in \lambda(t)\}$ is connected in T . We orient the tree from a chosen root vertex t_0 and enumerate the vertices of T as $T = \{t_i: i \leq n\}$ in such a manner that each set $T_i = \{t_j: j \leq i\}$ is connected. Let $u_j = \lambda(t_j) \in U_{s_j}$ so that $\mathcal{A} \upharpoonright u_j \simeq \mathcal{A}_{s_j}$ via a chart of \mathbb{A} , let $A_i = \bigcup_{j \leq i} u_j$, $\mathcal{A}_i = \mathcal{A} \upharpoonright A_i$. Let $\mathcal{C}_i = \mathcal{C} \upharpoonright A_i$, so that $\mathcal{C}_n = \mathcal{C}$. The desired homomorphism $h: \mathcal{C} \rightarrow \mathcal{B}$ is constructed from an increasing chain of homomorphisms $h_i: \mathcal{C}_i \rightarrow \mathcal{B}$ for $i = 1, \dots, n$ such that for each individual $j \leq i$ every $h_i \upharpoonright (C \cap u_j)$ extends to an isomorphism h_{ij}^* between $\mathcal{A} \upharpoonright u_j \simeq \mathcal{A}_{s_j}$ and some $\mathcal{B} \upharpoonright v_j$ for $v_j \in V_{s_j}$ that is induced by corresponding charts of \mathbb{A} and \mathbb{B} . Such h_i can be found inductively as follows. For $i = 0$, let $u_0 = \lambda(t_0) \in U_{s_0}$. We may then choose any $v_0 \in V_{s_0}$ in \mathbb{B} , and find an isomorphism between $\mathcal{A}_0 = \mathcal{A} \upharpoonright u_0 \simeq \mathcal{A}_{s_0}$ and $\mathcal{B} \upharpoonright v_0 \simeq \mathcal{A}_{s_0}$ as induced by corresponding charts of \mathbb{A} and \mathbb{B} ; the restriction of this isomorphism to $C \cap u_0$ can serve as h_0 . For the extension step from $h_i: \mathcal{A}_i \rightarrow \mathcal{B}$ to $h_{i+1}: \mathcal{A}_{i+1} \rightarrow \mathcal{B}$, let $t = t_{i+1}$, $s = s_{i+1}$ and $u = u_{i+1}$ be the building blocks associated with the new vertex t_{i+1} in T_{i+1} . Consider the immediate predecessor vertex t_j of t in the chosen orientation of T and let $h_{ij}^*(\mathcal{A} \upharpoonright u_j) = \mathcal{B} \upharpoonright v_j \simeq \mathcal{A}_{s_j}$ be the local extension of h_i to an isomorphism. Since \mathbb{B} also is a realisation of \mathbb{H}/\mathbb{I} , there is some $v \in V_s$ such that $\mathcal{B} \upharpoonright v \simeq \mathcal{A}_s$ intersects with $\mathcal{B} \upharpoonright v_j$ in $h_{ij}^*(u_j \cap u)$ in a manner that is compatible with the charts at u and v in \mathbb{A} and \mathbb{B} , respectively. To see this, we can use the fact that $u_j \cap u$ in \mathbb{A} is induced by some ρ_w for some walk w from s_j to s in \mathbb{I} ; and the succession of the same sequence of links from v_j in \mathbb{B} leads to a target v whose overlap with v_j must contain all of $h_{ij}^*(u_j \cap u)$. It follows that the extension of h_i by the isomorphism between $\mathcal{B} \upharpoonright v \simeq \mathcal{A}_s$ and $\mathcal{A} \upharpoonright u \simeq \mathcal{A}_s$, induced by corresponding charts, is well-defined in restriction to C , and meets the requirements for h_{i+1} . For well-definedness, as a homomorphism between relational structures, it is essential that the u_i form a tree decomposition of an atlas for \mathcal{C} : by Observation 4.8 there can be no identities between elements of \mathcal{C}_i and $\mathcal{C} \upharpoonright u$ other than those in $u_j \cap u$, and no relational links between any elements from $C \cap (u \setminus u_j)$ and from \mathcal{C}_i . \square

5 Two applications

5.1 Hypergraph coverings

As we saw in one of the first motivating examples, Example 2.3, any (finite) hypergraph (B, S) , with $S \subseteq \mathcal{P}_{\text{fin}}(B)$ as its set of hyperedges over the vertex set B , gives rise to an amalgamation pattern that describes an exploded view of (B, S) . This amalgamation pattern has (B, S) as a natural albeit trivial realisation. Recall from Definition 4.4 that we always assume $B = \bigcup S$. The associated incidence pattern is $\mathbb{I}(B, S)$ with S as its set of sites and $E = \{(s, s') \in S^2 : s \neq s', s \cap s' \neq \emptyset\}$ as the set of links. More formally, $\mathbb{I}(B, S) = (S, E, \iota, \cdot^{-1})$ is the 2-sorted encoding of the intersection graph of the set S of hyperedges, with directed edges and edge reversal. The amalgamation pattern that represents the exploded view of (B, S) over this incidence pattern $\mathbb{I} = \mathbb{I}(B, S)$ is

$$\mathbb{H}(B, S)/\mathbb{I} = (\mathbb{I}; H, \delta: H \rightarrow S; P, \eta: P \rightarrow E)$$

where

- $H = \bigcup_{s \in S} s \times \{s\}$ is the disjoint union of the hyperedges $s \in S$, partitioned by δ according to $\delta(a, s) = s$ so that $\delta^{-1}(s) = s \times \{s\} \subseteq H$ represents the hyperedge $s \in S$;
- P is the set of partial bijections $p_e = p_{(s, s')}: (s \cap s') \times \{s\} \rightarrow (s \cap s') \times \{s'\}$ representing the overlap between s and s' in (B, S) by mapping (b, s) to (b, s') for pairs $(s, s') = e \in E$ of hyperedges with a non-trivial intersection.

Clearly $\mathbb{H}(B, S)$ is strongly coherent and simple; in fact the action of \mathbb{I}^* through compositions ρ_w along walks in $\mathbb{I} = \mathbb{I}(B, S)$ is very simple: if $\iota(w) = (s, s')$, then ρ_w is the natural bijection between $(s \cap s') \times \{s\}$ and $(s \cap s') \times \{s'\}$. Also the symmetries of $\mathbb{H}(B, S)$ in the sense of Definition 2.7 are in bijective correspondence with the symmetries of (B, S) as a hypergraph, viz. with permutations of B that fix S as a set of subsets of B and thus also preserve overlaps. Obviously (B, S) itself can be cast as a realisation of $\mathbb{H}(B, S)/\mathbb{I}$, by associating each hyperedge $s \in S$ with the natural chart into $\delta^{-1}(s) = s \times \{s\}$ in \mathbb{H} . So (B, S) becomes the hypergraph of the co-ordinate domains of the atlas that connects (B, S) with its exploded view. In fact (B, S) corresponds to the quotient $\mathbb{H}(B, S)/\approx$ of the strongly coherent and simple $\mathbb{H}(B, S)/\mathbb{I}$ as in Observation 2.16. Proposition 5.3 below shows that, quite naturally, every realisation of $\mathbb{H}(B, S)/\mathbb{I}$ is a hypergraph covering of (B, S) in the sense of the following definition.

Definition 5.1. A hypergraph *homomorphism* between hypergraphs (A, U) and (B, S) is given by a map $\pi: A \rightarrow B$ whose restriction to each hyperedge $u \in U$ bijectively maps that u onto an image hyperedge $\pi(u) \in S$. We naturally extend π to the second sort (of hyperedges) and denote homomorphisms as in $\pi: (A, U) \rightarrow (B, S)$.

A homomorphism $\pi: (A, U) \rightarrow (B, S)$ is a *covering* of the hypergraph (B, S) by the hypergraph (A, U) if it maps U surjectively onto S (and hence also A onto B), and satisfies the following lifting property for overlaps between hyperedges:

for every pair of hyperedges (s, s') of (B, S) with nontrivial overlap, and every $u \in U$ with $\pi(u) = s$ there is some $u' \in U$ s.t. $\pi(u') = s'$ and $s \cap s' = \pi(u \cap u')$ (equivalently: s.t. π restricts to a bijection of $u \cup u'$ onto $s \cup s'$).

A covering $\pi: (A, U) \rightarrow (B, S)$ is *fully symmetric* over (B, S) if every symmetry of (B, S) (or of $\mathbb{H}(B, S)$) extends to a symmetry of $\pi: (A, U) \rightarrow (B, S)$ and if the subgroup of B -rigid symmetries (symmetries that fix B , and therefore $\mathbb{H}(B, S)$, pointwise) acts transitively on $\pi^{-1}(s) \subseteq U$ for every $s \in S$.

Note that our notion of a covering is what in other contexts is called a *branched covering*, since the multiplicities of overlaps between hyperedges in $\pi^{-1}(s) \subseteq U$ and $\pi^{-1}(s') \subseteq U$ may be greater than one. Simple examples show that coverings of a given (finite) hypergraph by (finite or infinite) hypergraphs that are acyclic, or even just N -acyclic for specific N , may be necessarily branched. (The situation for graphs is quite different: finite unbranched coverings by N -acyclic graphs are always available, and can be obtained as direct products with finite groups of large girth [18].)

Example 5.2. Consider the hypergraph with 4 vertices and 3 hyperedges of size 3 that corresponds to a triangulation of a triangle with one extra central vertex. This hypergraph is not 3-acyclic, and every non-trivial covering hypergraph must have higher multiplicities at the central vertex (infinite for a properly acyclic covering).

Proposition 5.3. *Every realisation \mathbb{A} of $\mathbb{H}(B, S)/\mathbb{I}$ for $\mathbb{I} = \mathbb{I}(B, S)$ induces a covering of the hypergraph (B, S) by the hypergraph (A, U) associated with the atlas of the realisation \mathbb{A} . This covering is fully symmetric over (B, S) if the realisation \mathbb{A} is fully symmetric over $\mathbb{H}(B, S)$, and the covering hypergraph is N -acyclic if the realisation \mathbb{A} is N -acyclic.*

Proof. Let \mathbb{A} be a realisation of $\mathbb{H}(B, S)/\mathbb{I}$ with atlas $(A, U, (U_s), (\pi_{u,s}))$ so that every $\pi_{u,s}$ for $u \in U_s$ bijectively maps $u_s \subseteq A$ to $s \times \{s\} \subseteq H$. Let $(\pi_{u,s})_1$ be the composition of $\pi_{u,s}$ with projection to the first component of the image, which is a bijection from $u \in U_s$ onto s . If $a \in u \cap u'$ for $u \in U_s$ and $u' \in U_{s'}$, then the properties of \mathbb{A} as a realisation imply that, for some walk w from s to s' in \mathbb{I} , $\pi_{u',s'}(a) = \rho_w(\pi_{u,s}(a))$ in $\mathbb{H}(B, S)$, which implies that $(\pi_{u',s'})_1(a) = (\pi_{u,s})_1(a)$ by the special nature of the ρ_w in $\mathbb{H}(B, S)$. It follows that the global union of all the $(\pi_{u,s})_1$ is a well-defined surjection $\pi: A \rightarrow B$ of A onto $B = \bigcup S$. It remains to show that $\pi: (A, U) \rightarrow (B, S)$ is a hypergraph covering; the claims regarding symmetry and N -acyclicity are then obvious. Clearly π is a hypergraph homomorphism that is surjective also at the level of hyperedges. The lifting property for non-trivial overlaps between hyperedges $s, s' \in S$ at some $u \in U_s$ is obvious from the corresponding condition on realisations: since $s \cap s' \neq \emptyset$ and $s \neq s'$, $e = (s, s')$ is a link in $\mathbb{I} = \mathbb{I}(B, S)$, which has to be realised in \mathbb{A} by the overlap with some $u' \in U_{s'}$, and compatibility with ρ_e (in the sense of condition (i) of Definition 2.14) means that $u \cap u'$ precisely matches $s \cap s'$. \square

Corollary 5.4. *For every N , every finite hypergraph admits fully symmetric finite coverings by N -acyclic hypergraphs.*

Example 5.5. Consider 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. Its Gaifman graph does not have any chordless cycles of length greater than 3, but conformality is violated by the 4-clique of all its vertices. A finite N -acyclic covering for $N \geq 4$ cannot be unbranched since it must unravel the local 3-cycles of hyperedges that meet in a single vertex into a cycles of lengths $3n$ for some $n > 1$. While it is not hard to visualize a natural fully symmetric covering based on a quotient of a hexagonal grid pattern (locally two-fold, for $n = 2$), which is 5-acyclic, there seems to be no obvious ad-hoc construction for higher degrees of acyclicity as guaranteed by Corollary 5.4.

5.2 Lifting local to global symmetries

Recall that a partial isomorphism $p \in \text{part}(\mathcal{A}, \mathcal{A}')$ is an isomorphism between induced substructures of \mathcal{A} and \mathcal{A}' . In case $\mathcal{A} = \mathcal{A}'$, we speak *partial automorphisms*, which are partial or local symmetries within structure \mathcal{A} , i.e. elements of the inverse semigroup of partial bijections that respect the interpretations of given relations of \mathcal{A} . It is not hard to see that a collection of partial automorphisms in $\text{part}(\mathcal{A}, \mathcal{A})$ can simultaneously be extended to global automorphisms in the automorphism group $\text{aut}(\mathcal{B})$ of some infinite extension $\mathcal{B} \supseteq \mathcal{A}$. A straightforward construction of such solutions to the *extension problem for partial automorphisms* could proceed along the lines of Remark 4.3 for the construction of infinite realisations as reduced products. The problem becomes more subtle if, for finite \mathcal{A} and a collection of its partial automorphisms, one asks for *finite* solutions or for *finite* solutions within a restricted class of structures. Another important variant of the problem asks for *finite* solutions within a restricted class \mathcal{C} provided that at least infinite solutions exist in \mathcal{C} . Note the conditional character of this variant, which is in the style of a finite model property for the extension task. Classes \mathcal{C} of relational structure that satisfy this version are said to have the *extension property for partial automorphisms* (for short: EPPA) in terminology introduced in [13].

Classical results for the unconditional extension task in restriction to finite structures were first obtained for the class of all finite graphs by Hrushovski in [15], then for the class of all finite structures in any fixed finite relational signature by Herwig in [11]. The strongest and most general result, in the direction of the conditional (EPPA) for restricted classes of structures, is the result of Herwig and Lascar in [13], Theorem 5.8 below: any class of relational structures defined in terms of finitely many finite *forbidden homomorphisms* satisfies (EPPA). This result subsumes the earlier, unconditional results for classes of finite structures, since infinite solutions are always available in those cases. For the result for finite graphs, a very elegant, elementary and direct proof has been presented in [13], which avoids both the substantial intricacies of the much stronger general (EPPA) result in [13] and the non-trivial algebraic arguments of the original proof in [15].

As it turns out, fully symmetric finite realisations of naturally induced amal-

gamation patterns offer an alternative direct route to particularly generic finite solutions for the extension problem for finite automorphisms. This connection is especially striking in the general case of the conditional Herwig–Lascar result, where the insights from Section 4.5 into the connection between local acyclicity and local genericity up to homomorphisms play out in a particularly nice way.

Definition 5.6. Let σ be a finite relational signature, \mathcal{C} a class of σ -structures. An *instance* of the *extension problem for partial automorphisms* over \mathcal{C} is a pair (\mathcal{A}_0, P) consisting of a finite σ -structure $\mathcal{A}_0 \in \mathcal{C}$ and a set P of partial isomorphisms in \mathcal{A}_0 (partial automorphisms, local symmetries of \mathcal{A}_0); a *solution* for this instance is any \mathcal{A} in \mathcal{C} into which \mathcal{A}_0 embeds isomorphically such that the image of each $p \in P$ extends to a full automorphism of \mathcal{A} .

The class \mathcal{C} has the *extension property for partial automorphisms* (EPPA) if every instance (\mathcal{A}_0, P) of the extension problem over \mathcal{C} that admits any solution also admits a finite solution in \mathcal{C} .

In the following we shall always assume that the set P of partial automorphisms of an instance is closed under inverses: for each $p \in P$, also $p^{-1} \in P$. This is w.l.o.g. in the sense that p^{-1} can be adjoined to P without changing the set of solutions.

For classes \mathcal{C} like the class of all graphs or of all σ -structures, the condition that an instance has a solution in \mathcal{C} is universally fulfilled, so that (EPPA) for such classes boils down to the requirement that every instance over \mathcal{C}_{fin} admits a solution in \mathcal{C}_{fin} , where \mathcal{C}_{fin} stands for the subclass of finite structures in \mathcal{C} .

Definition 5.7. A class \mathcal{C} of σ -structures is defined in terms of *forbidden homomorphisms* if there is some finite set X of finite σ -structures such that \mathcal{C} is the class of those σ -structures \mathcal{A} that admit *no* homomorphism $h: \mathcal{X} \rightarrow \mathcal{A}$ from any $\mathcal{X} \in X$.

Theorem 5.8 (Herwig–Lascar). *Any class \mathcal{C} of σ -structures that is defined in terms of forbidden homomorphisms has the extension property for partial automorphisms (EPPA).*

Positive examples include the classes of all (finite) graphs [15], of (finite) K_n -free graphs [13], of all (finite) σ -structures [11]. Known examples of classes \mathcal{C} with (EPPA) that are not defined in terms of forbidden homomorphisms include the class of (finite) conformal σ -structures, i.e. structures for which all cliques in the Gaifman graph are generated by single relational atoms [14]. It appears to be still open whether the class of all tournaments or the class of all 4-cycle-free graphs have (EPPA).

Definition 5.9. We call a solution \mathcal{A} to the instance (\mathcal{A}_0, P) of the extension problem a *fully symmetric solution* if it possesses an atlas $(\mathcal{A}, U, (\pi_u)_{u \in U})$ of co-ordinate maps $\pi_u: \mathcal{A} \upharpoonright u \simeq \mathcal{A}_0$ onto the single co-ordinate structure \mathcal{A}_0 such that

- (i) the automorphism group of \mathcal{A} acts transitively on the set U of the co-ordinate domains in a manner compatible with the charts (i.e. by symmetries of the atlas that are rigid w.r.t. the single co-ordinate structure);

- (ii) any two co-ordinate domains $u, u' \in U$ with non-trivial intersection are related by the action of a composition of partial isomorphisms from P in the sense that $p_{u'} \circ p_u^{-1} = p_n \circ \dots \circ p_1$ for a sequence $w = p_1 \dots p_n \in P^*$.¹²

Theorem 5.10. *For the class \mathcal{C} of all σ -structures and any $N \in \mathbb{N}$: every instance (\mathcal{A}_0, P) for $\mathcal{A}_0 \in \mathcal{C}$ admits finite solutions $\mathcal{A} \in \mathcal{C}_{\text{fin}}$ which are fully symmetric and whose atlas is N -acyclic.*

Proof. The desired finite solution is obtained from a finite, fully symmetric and N -acyclic realisation of an amalgamation pattern $\mathbb{H}(\mathcal{A}_0, P)$ derived from the extension task (\mathcal{A}_0, P) as follows. We use the incidence pattern $\mathbb{I} = \mathbb{I}(\mathcal{A}_0, P) = \mathbb{I}_P$ with the singleton set $S = \{0\}$ for its sort of sites, and edges $e_p \in E$ for each $p \in P$, with constant $\iota(p) = (0, 0)$ and with edge reversal according to $e_p^{-1} = e_{p^{-1}}$. We may identify the set P with the sort E of links of \mathbb{I}_P . Formally, the amalgamation pattern $\mathbb{H}(\mathcal{A}_0, P)$ then is

$$\mathbb{H}(\mathcal{A}_0, P)/\mathbb{I}_P = (\mathbb{I}; \mathcal{A}_0, 0: A_0 \rightarrow \{0\}; P, \text{id}: P \rightarrow P)$$

with the constant function 0 (as δ) for the trivial S -partition of $H = A_0$, and the identity function on P (for η) that labels the partial isomorphism $p \in \text{part}(\mathcal{A}_0, \mathcal{A}_0)$ by its name, which is the link $p = e_p \in E = P$.

We claim that every realisation $\mathbb{A}/\mathbb{H} = (\mathbb{H}, \mathcal{A}, U, (\pi_u)_{u \in U})$ of $\mathbb{H} = \mathbb{H}(\mathcal{A}_0, P)$ that is fully symmetric over $\mathbb{H}(\mathcal{A}_0, P)$ is a solution to the extension problem for (\mathcal{A}_0, P) . More specifically, the σ -structure \mathcal{A} of \mathbb{A} embeds an isomorphic copy of \mathcal{A}_0 as a substructure in such a way that the image of every $p \in \text{part}(\mathcal{A}_0, \mathcal{A}_0)$ extends to some automorphism of \mathcal{A} .

Any embedding of \mathcal{A}_0 according to one of the charts $\pi_u: \mathcal{A} \upharpoonright u \simeq \mathcal{A}_0$ can be used. We fix some $u_0 \in U$ and regard $\mathcal{A} \upharpoonright u_0$ as the distinguished copy of \mathcal{A}_0 in \mathcal{A} . Let its chart be $\pi_0 := \pi_{u_0}: \mathcal{A} \upharpoonright u_0 \simeq \mathcal{A}_0$. The incarnation of $p \in \text{part}(\mathcal{A}_0, \mathcal{A}_0)$ in this distinguished copy of \mathcal{A}_0 is $\pi_0^{-1} \circ p \circ \pi_0 \in \text{part}(\mathcal{A} \upharpoonright u_0, \mathcal{A} \upharpoonright u_0)$.

For $p \in P$ choose $u \in U$ for e_p as in condition (i) of Definition 2.14 for realisations, such that

$$\pi_u \circ \pi_0^{-1} = p \in \text{part}(\mathcal{A}_0, \mathcal{A}_0), \quad (1)$$

where the composition on the left is a partial composition of the restrictions that fit together precisely in $\pi_0^{-1}(\text{dom}(p)) = \pi_u^{-1}(\text{image}(p))$.

As \mathbb{A} is fully symmetric over \mathbb{H} , it possesses an \mathbb{H} -rigid symmetry π that maps u to u_0 in a manner that is compatible with π_u and π_0 :

$$\pi_0 \circ \pi = \pi_u, \quad (2)$$

where the composition on the left is defined on all of u .

¹²The equality of partial compositions includes the equality of the domains. Recall that we assume P to be closed under inverses; P^* can be understood as either the set of finite words over the alphabet P , or, equivalently as the set of walks in an incidence pattern $\mathbb{I}_P = (\{0\}, P)$ with a site 0 and loops for $p \in P$.

Let $a \in u_0$ be in the domain of $\pi_0^{-1} \circ p \circ \pi_0$ (i.e. p in the distinguished copy of \mathcal{A}_0 in \mathcal{A}). Then $p(\pi_0(a)) = \pi_u(a)$ by (1), so that $\pi_0^{-1}(p(\pi_0(a))) = \pi_0^{-1}(\pi_u(a)) = \pi(a)$ by (2). So the symmetry π of \mathbb{A} extends $\pi_0^{-1} \circ p \circ \pi_0 \in \text{part}(\mathcal{A} \upharpoonright u_0, \mathcal{A} \upharpoonright u_0)$ to an automorphism of \mathcal{A} .

The extra symmetry and acyclicity properties of the solution \mathcal{A} with its natural atlas are immediate from corresponding properties of \mathbb{A} . \square

Observation 5.11. *Considering the manner in which suitable finite realisations of \mathbb{H}/\mathbb{I} were obtained as reduced products with suitable groupoids \mathbb{G}/\mathbb{I} in Section 4, it is worth noting that the incidence pattern \mathbb{I}_P for $\mathbb{H}(\mathcal{A}_0, P)/\mathbb{I}_P$ is such that any groupoid \mathbb{G}/\mathbb{I}_P is in fact a Cayley group.*

The following strengthening of the Herwig–Lascar result of Theorem 5.8 is a consequence of Theorem 5.10 and our considerations about N -acyclicity and N -universality in Section 4.5. In analogy with Definition 4.13 we say that a σ -structure \mathcal{A} is N -universal within a class \mathcal{C}_0 of σ -structures (in our case, a class of solutions to an instance of an extension problem) if every substructure of \mathcal{A} of size up to N admits homomorphisms to every structure in \mathcal{C}_0 .

Corollary 5.12. *Consider the instance (\mathcal{A}_0, P) of the extension problem in the class of all σ -structures and let \mathcal{A} be any finite, N -acyclic and fully symmetric solution as in Theorem 5.10.*

- (a) *Then \mathcal{A} is N -universal within the class $\mathcal{C}_0 = \mathcal{C}_0(\mathcal{A}_0, P)$ of all solutions to the instance (\mathcal{A}_0, P) .*
- (b) *Let the class \mathcal{C} be defined in terms of forbidden homomorphisms and N sufficiently large, viz. an upper bound for the sizes of the forbidden homomorphisms that define \mathcal{C} . Then $\mathcal{A} \in \mathcal{C}_{\text{fin}}$ if the instance (\mathcal{A}_0, P) has any finite or infinite solution in \mathcal{C} .*

Proof. Claim (b) — which in particular implies (EPPA) for the class \mathcal{C} — is an immediate consequence of (a). We deal with the more specific claim (a).

Let \mathcal{A} be a finite N -acyclic, fully symmetric solution as in the proof of Theorem 5.10 with the atlas $(\mathcal{A}, U, (U_s), (\pi_u)_{u \in U})$ of charts $\pi_u: \mathcal{A} \upharpoonright u \simeq \mathcal{A}_0$. Let \mathcal{B} be any other solution for the instance (\mathcal{A}_0, P) of the extension problem over the class of all σ -structures, with $\mathcal{A}_0 \simeq \mathcal{B}_0 \subseteq \mathcal{B}$ for the embedded copy of \mathcal{A}_0 . For $p \in P$, let $\pi_p \in \text{Aut}(\mathcal{B})$ be an extension of the image of $p \in \text{part}(\mathcal{A}_0, \mathcal{A}_0)$ in $\mathcal{B}_0 \subseteq \mathcal{B}$. We need to construct, for a given $\mathcal{D} \subseteq \mathcal{A}$ of size up to N , a homomorphism $h: \mathcal{D} \rightarrow \mathcal{B}$.

As in the proof of Proposition 4.14, we can use a tree decomposition of \mathcal{D} that is induced by the N -acyclic atlas of \mathcal{A} . This tree decomposition will govern an inductive choice of a sequence of homomorphisms h_i for substructures $\mathcal{D}_i \subseteq \mathcal{D}$. By N -acyclicity of the atlas, $\mathcal{D} = \mathcal{A} \upharpoonright \mathcal{D} \subseteq \mathcal{A}$ for $|\mathcal{D}| \leq N$, \mathcal{D} admits a tree decomposition $(T, E, \lambda: T \rightarrow U)$ such that $\mathcal{D} \subseteq \bigcup_{t \in T} \lambda(t)$ and such that the sets $\{t \in T: d \in \lambda(t)\} \subseteq T$ are connected for all $d \in \mathcal{D}$. Let t_0 be the root of T w.r.t. a chosen orientation and enumerate T as $T = \{t_i: i \leq n\}$ such all initial subsets $T_i = \{t_j: j \leq i\} \subseteq T$ are connected. Let $\lambda(t_j) = u_j \in U$ so that $\pi_j := \pi_{u_j}: \mathcal{A} \upharpoonright u_j \simeq \mathcal{A}_0$. Let $A_i = \bigcup_{j \leq i} u_j$, $\mathcal{A}_i = \mathcal{A} \upharpoonright A_i$, and $\mathcal{D}_i = \mathcal{D} \upharpoonright A_i$.

The homomorphism $h: \mathcal{D} \rightarrow \mathcal{B}$ is the final member of an increasing sequence of homomorphisms $h_i: \mathcal{D}_i \rightarrow \mathcal{B}$ for $i = 1, \dots, n$ with this additional property: for each $j \leq i$, $h_i \upharpoonright (D \cap u_j)$ extends to an isomorphism $h_{ij}^*: \mathcal{A} \upharpoonright u_j \simeq \mathcal{B}_j \subseteq \mathcal{B}$ where $\mathcal{B}_j \simeq \mathcal{A}_0$ is an image of $\mathcal{A}_0 \simeq \mathcal{B}_0 \subseteq \mathcal{B}$ under some composition π_{w_j} of the designated extensions π_p of the $p \in P$ to automorphisms of \mathcal{B} , corresponding to some word $w_j = p_1 \cdots p_n \in P^*$, according to

$$\mathcal{B}_j = \pi_{w_j}(\mathcal{A}_0) = (\pi_{p_n} \circ \cdots \circ \pi_{p_1})(\mathcal{A}_0) \subseteq \mathcal{B}.$$

Such h_i can be chosen inductively as follows. For $i = 0$, let $u_0 = \lambda(t_0) \in U$. We may combine the isomorphic embedding of $\mathcal{A}_0 \simeq \mathcal{B}_0 \subseteq \mathbb{B}$ and the isomorphism between $\pi_{u_0}: \mathcal{A} \upharpoonright u_0 \simeq \mathcal{A}_0$ to obtain the desired isomorphism h_{00}^* , and use its restriction to $D \cap u_0$ as h_0 .

For the extension step from $h_i: \mathcal{A}_i \rightarrow \mathcal{B}$ to $h_{i+1}: \mathcal{A}_{i+1} \rightarrow \mathcal{B}$, let $u = \lambda(t_{i+1})$ and consider the immediate predecessor t_j of t_{i+1} in T . Let $h_{ij}^*(\mathcal{A} \upharpoonright u_j) = \mathcal{B}_j \simeq \mathcal{A}_0$ be the local extension of h_i to an isomorphism, where $\mathcal{B}_j = \pi_{w_j}(\mathcal{A}_0) \subseteq \mathcal{B}$.

Since \mathcal{A} is a fully symmetric solution, $u_j \cap u$ in \mathbb{A} is induced by some $\rho_{w'}$ for a word $w' \in P^*$. The corresponding operation of the composition $\pi_{w'}$ acts as an automorphism on \mathcal{B} and thus produces an overlap between $\mathcal{B}_j = \pi_{w_j}(\mathcal{A}_0)$ and $\mathcal{B}_{i+1} := \pi_{w_j w'}(\mathcal{A}_0) = \pi_{w'}(\mathcal{B}_j)$ that contains the h_{ij}^* -image of $\mathcal{A} \upharpoonright (u_j \cap u)$. We can use this \mathcal{B}_{i+1} , the natural isomorphism induced by the chart of $\mathcal{A} \upharpoonright u$ in the atlas of \mathcal{A} and the automorphism $\pi_{w_j w'}^{-1}$ that maps \mathcal{B}_{i+1} onto $\mathcal{B}_0 \subseteq \mathcal{B}$ as $h_{i+1, i+1}^*$ and join its restriction to $u \cap D$ to h_i in order to obtain h_{i+1} . The argument for well-definedness as a homomorphism between relational structures is strictly analogous to that in the proof of Proposition 4.14, based on the properties of a tree decomposition. \square

References

- [1] S. Abramsky. Relational hidden variables and non-locality. *Studia Logica*, 101(2), pp. 411–452, 2013.
- [2] 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.
- [3] 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*.
- [4] 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*.
- [5] 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.
- [6] C. Berge. *Graphs and Hypergraphs*. North-Holland, 1973.

- [7] 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.
- [8] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *Annals of Pure and Applied Logic*, 161: pp. 1–42, 2009.
- [9] R. Fox. Covering spaces with singularities. In *A Symposium in Honor of S. Lefschetz*, pp. 243–257. Princeton University Press, 1957.
- [10] 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.
- [11] B. Herwig. Extending partial isomorphisms on finite structures. *Combinatorica*, 15, pp. 365–371, 1995.
- [12] B. Herwig. Extending partial isomorphisms for the small index property of many omega-categorical structures. *Israel Journal of Mathematics*, 107, pp. 93–124, 1998.
- [13] B. Herwig and D. Lascar. Extending partial isomorphisms and the profinite topology on free groups. *Transactions of the AMS*, 352, pp. 1985–2021, 2000.
- [14] I. Hodkinson and M. Otto. Finite conformal hypergraph covers and Gaifman cliques in finite structures. *Bulletin of Symbolic Logic*, 9, pp. 387–405, 2003.
- [15] E. Hrushovski. Extending partial isomorphisms of graphs. *Combinatorica*, 12, pp. 411–416, 1992.
- [16] M. Lawson. *Inverse Semigroups*. World Scientific, 1998.
- [17] D. Macpherson. A survey of homogeneous structures. *Discrete Mathematics*, 311, pp. 1599–1634, 2011.
- [18] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130: pp. 173–205, 2004.
- [19] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59 (1), 2012.
- [20] M. Otto. On groupoids and hypergraphs. Technical report, 2012. See *arXiv:1211.5656*.
- [21] M. Otto. Groupoids, hypergraphs and symmetries in finite models. In *Proceedings of 28th Annual IEEE Symposium on Logic in Computer Science LICS'13*, 2013. Also see *arXiv:1404.4599*, 2016.