Highly Acyclic Groups, Hypergraph Covers and the Guarded Fragment

Martin Otto
Department of Mathematics
Technische Universität Darmstadt

Abstract

We construct finite groups whose Cayley graphs have large girth even w.r.t. a discounted distance measure that contracts arbitrarily long sequences of edges from the same colour class, and only counts transitions between colour classes. These groups are shown to be useful in the construction of finite bisimilar hypergraph covers that avoid any small cyclic configurations. We present two applications to the finite model theory of the guarded fragment: a strengthening of the known finite model property for GF *and the characterisation of* GF *as the guarded bisimulation invariant fragment of* FO *in the sense of finite model theory.*

1. Introduction

Acyclicity of hypergraphs [5] or relational structures has long been recognised as an important structural property because of its relation with tree decomposability [4]. Acyclicity criteria or criteria for tree decomposability, often also in the more liberal form of bounds on tree width (and generalisations), often play an important role in the delineation of well-behaved problem instances, e.g., for model checking or query answering. But also from a purely model theoretic point of view tree decomposable models – and again, more liberally, models of bounded tree width – are of interest because of their interpretability in actual trees, which makes them amenable, for instance, to automata theoretic techniques. Erich Grädel's *generalised tree model property* for the guarded fragment GF $\subseteq$ FO in [7], which is responsible for a range of decidability results and complexity bounds, is an important case in point. The natural notion of unfolding of relational structures, which is compatible with guarded bisimulation equivalence, produces models that are tree-like not just in the sense of bounded tree width, but of tree decomposability of the hypergraph of guarded subsets of the model. Similar phenomena are well known from graph-like structures (especially transition systems), which can be unfolded into bisimilar tree structures.

If we want to stick with finite structures, then these tree unfoldings are not available since, even in the graph case, any cycle in the original structure can only be unfolded into an infinite path. For some purposes, however, it suffices to achieve some measure of local or bounded acyclicity rather than global acyclicity. In the case of graph-like structures, constructions of bisimilar covers by finite, *locally acyclic* structures are available [10] and have been used for constructive alternative proofs of expressive completeness results for modal logics, which, unlike classical proof methods, work in the context of finite model theory and other non-elementary classes of structures [6].

The situation for hypergraphs or relational structures of width greater than 2, as opposed to graphs or graph-like structures, has proved a major challenge in this respect. While it is quite clear what the natural notion of bisimilar *hypergraph covers* ought to be, is not at all obvious which measure of hypergraph acyclicity can be achieved in finite covers. The immediate analogue of the graph case is ruled out: *local acyclicity* – in the sense that the hypergraph structures induced on ℓ-neighbourhoods must be acyclic – cannot be achieved. What, then, can be hoped for?

Classical hypergraph theory [5, 4] characterises acyclicity and tree-decomposability of hypergraphs in terms of two independent requirements: *conformality* and *chordality*. The former forbids cliques in the Gaifman graph other than those that are induced by individual hyperedges; the latter forbids chordless cycles in the Gaifman graph.

Conformal finite covers were constructed in [9] and employed in a simplified proof of the finite model property for the clique guarded fragment.

Chordal finite covers cannot in general be had. The simplest counterexample is a cartwheel hypergraph consisting of at least three 3-hyperedges that all share one pivot vertex and form a cycle w.r.t. the edges formed by the remaining two vertices in each hyperedge. This same example rules out local chordality even in 1-neighbourhoods.

A natural relaxation would forbid just *short* chordless cycles (of lengths up to N say), which we call N*-chordality*.

A construction of finite hypergraph covers in which short cycles in the cover admit a chordal decomposition in projection to the base (weak N-chordality) was obtained in [11]; an alternative, much more explicit construction of such covers with reasonable size bounds is now presented in [3].

These approaches notwithstanding, the question whether finite hypergraphs generally admit finite conformal and N-chordal covers remained wide open; and a positive solution a likely indispensable ingredient in any expressive completeness proof for GF in finite models.

With methods entirely different from those in [11] or [3] we here now obtain conformal and N-chordal finite hypergraph covers. Our main theorem in this respect is Theorem 16, which in turn is based on the combinatorial main result, about highly acyclic Cayley groups, presented as Corollary 8. Both results are technically rather involved.

These results point us to the class of relational structures of *bounded acyclicity*, which seems to be very smooth from a model theoretic point of view. For instance, it supports a natural notion of *bounded convex hulls*. These insights are used to settle one of the key open questions about the finite model theory of the guarded fragment. We show that the guarded fragment is expressively complete for all first-order properties that are invariant under guarded bisimulation in finite models, Theorem 26. The proof uses conformal N-chordal covers and further applications of our main combinatorial result, plus the analysis of bounded acyclicity. It gives an essentially constructive expressive completeness argument totally different from the classical variant due to Andréka, van Benthem and Németi, [2], and shows the guarded fragment to behave in beautiful analogy with the modal fragment in yet another way.

The paper is organised in three separate parts, each starting with a short review of the relevant technical notions. The first, Section 2, culminates in Corollary 8; it deals with the combinatorial constructions that are the main technical tool in the following; the Cayley groups obtained may well be of independent interest. The second part, Section 3, develops the construction of conformal N-chordal hypergraph covers as stated in Theorem 16. The third part, Section 4, deals with the analysis of bounded acyclicity in connection with GF; as an immediate application we prove a strengthened finite model property for GF in Corollary 18; substantially more work is required to prove the expressive completeness result for GF in Theorem 26.

2. Highly acyclic groups

We aim to construct finite, regularly edge-coloured, homogeneous graphs which do not realise short cycles, or even cycles that would be short when subjected to certain contractions of paths running within the same group of colours.

2.1. Cayley groups and graphs

Regular graphs of large girth A simple example is the following construction [1] of k-regular graphs of girth greater than N, for arbitrary given k and N. Let T be the regularly k-coloured undirected tree, in which every node has precisely one neighbour across each one of the k edge-colours $e_1, \dots, e_k$; designate one node λ in this tree as its root and truncate the whole tree at depth N from the root. Each colour e_i now induces a permutation π_i of the vertex set of this finite tree, if we let π_i swap each pair of vertices that are linked by an e_i-edge (note that just some leaves are fixed by π_i). Let G be the group generated by $(\pi_i)_{1 \leqslant i \leqslant k}$ in the full symmetric group of the vertex set. We obtain the desired graph as the Cayley graph of the group G: its vertices are the group elements $g \in G$; g and g' are linked by an edge (of colour e_i) if $g' = g \circ \pi_i$ (equivalently: $g = g' \circ \pi_i$, as π_i is involutive). It is clear that no non-trivial sequence of generators of length up to N can represent the neutral element $1 \in G$: just observe its operation on the root λ to see that λ is moved precisely one step away from the root by each new generator application, whence a sequence of up to N generators cannot be trivial. It follows that the Cayley graph has girth greater than N.

In the following we shall modify the basic idea in this construction to yield finite graphs displaying a much stronger form of acyclicity w.r.t. to discounted distance measures along cycles. For the rest of this section, let E be a finite set of *edge colours*. A subset $\alpha \subseteq E$ is regarded as a *colour class*. We deal with E-coloured undirected graphs in which every node is incident with at most one edge of any fixed colour e. We call such graphs *E-graphs*. The class of E-graphs is closed under subgraphs (also in the sense of weak substructures) as well as under reducts.

In E-graphs, connected components w.r.t. subsets $\alpha \subseteq E$ are defined as usual. We here regard an α-component of an E-graph as an α-graph, in the sense of an implicit passage to the α-reduct, with all edges of colours $e \notin E$ removed. We shall in particular look at Cayley graphs of groups generated by a set of pairwise distinct involutive generators $e \in E$; in the following we just speak of generators $e \in E$ and *groups with generator set E*. In any such group G we associate with the word $w = e_1 \cdots e_n$ over E the group element $[w]^G = e_1 \circ \cdots \circ e_n$. We think of the letters e_i also as edge labels along a path w from 1 to $[w]^G$ in the Cayley graph of G; in the natural fashion we let G operate on its Cayley graph from the right, so that $e_i = [e_i]^G$ translates g into $g \circ e_i$. We denote by w^{-1} the word $w^{-1} = e_n \cdots e_1$ obtained by reversing $w = e_1 \cdots e_n$; clearly $[w^{-1}]^G = ([w]^G)^{-1}$ because of the involutive nature of the generators.

For any such group G we also denote its Cayley graph by G, which is a regular E-graph. For a subset $\alpha \subseteq E$ we look at the subgroup $G_\alpha := G \restriction \alpha \subseteq G$ generated by this subset; its Cayley graph is naturally isomorphic to the α-component of 1 in the Cayley graph of G, which is an α-graph (as well as an E-graph).

If H is any E-graph, we write $\mathrm{sym}(H)$ for the Cayley group induced by the natural operation of edge colours $e \in E$ as involutions, as reviewed above. (The lower case $\mathrm{sym}(H)$ distinguishes this subgroup from the full symmetric group of the set H.) If G is the Cayley graph of a group G with generator set E, then the group G is reproduced as $\mathrm{sym}(G)$.

Definition 1. Let G be a group with generator set E.

(i) An E-graph H is *compatible* with G if for all words w over E: $[w]^G = 1 \;\Rightarrow\; [w]^{\mathrm{sym}(H)} = 1$.

(ii) G *reflects intersections* if, for all $\alpha, \beta \subseteq E$: $G_\alpha \cap G_\beta = G_{\alpha\cap\beta}$.

If $G = \mathrm{sym}(H)$, then every connected component of H is trivially compatible with G. If H' is compatible with $G = \mathrm{sym}(H)$, then $G \simeq \mathrm{sym}(H \dot{\cup} H')$; i.e., G is unaffected by adjoining any compatible graph as a disjoint component.

It is not hard to see that G reflects intersections if (the Cayley graphs of) its subgroups G_α are compatible with G, for all $\alpha \subseteq E$.

Remark 2. *Suppose G_α is compatible with G for every $\alpha \subseteq E$. Then $G = \mathrm{sym}(G) = \mathrm{sym}(G \cup \dot{\bigcup}_\alpha G_\alpha)$ and for any $h = \prod_i e_i$ and $\alpha \subseteq E$:*
$h \in G_\alpha \Rightarrow h = h \restriction \alpha := \prod_{i\colon e_i \in \alpha} e_i$, with all $e_i \notin \alpha$ deleted. Hence G reflects intersections.

Using the same notion of $h \restriction \alpha$, it is useful to observe that G_α is compatible with G iff, for all $h = \prod_i e_i$,

$$h = 1 \text{ (in } G) \quad \Rightarrow \quad h \restriction \alpha = 1 \text{ (in } G_\alpha);$$

in this situation $h \restriction \alpha$ is well defined in terms of h (rather than in terms of a particular representation $h = \prod_i e_i$).

Merging chains of components Consider any two E-graphs K and K' with distinguished nodes $v \in K, v' \in K'$ and a distinguished subset $\alpha \subseteq E$. Assume that the α-components of v and v' are isomorphic via some isomorphism ρ that maps v to v'. We let

$$K \;\frac{v=v'}{\alpha}\; K'$$

be the result of glueing K and K' according to the isomorphism ρ. If α is the intersection of the colour classes of K and K', then this merged graph is again an E-graph.

In the following we shall build chains by merging α-components of the Cayley graph of a group G. In this case there always is, for any nodes $g \in G_\alpha$ and $g' \in G_{\alpha'}$, a unique isomorphism between the $(\alpha \cap \alpha')$-components of $g \in G_\alpha$ and of $g' \in G_{\alpha'}$ that maps g to g'.

In merging a sequence of graphs $(K_s)_{1 \leqslant s \leqslant n}$, each with designated nodes to be identified with corresponding nodes in the left and right neighbours, we perform these identifications simultaneously, i.e., apply the isomorphisms between matching components in any pair of neighbours along the sequence. A simple sufficient condition that guarantees that the resulting graph is again an E-graph, is the following: we require the two patches in K_s that are joined with patches in K_{s-1} and K_{s+1}, respectively, to be disjoint; in this manner no identifications are carried through any three or more consecutive members in the merged chain.

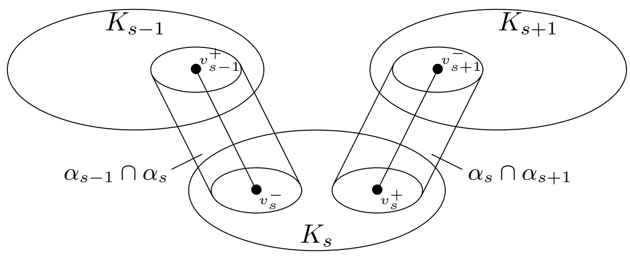

Definition 3. Consider a sequence $(K_s, v_s^-, v_s^+)_{1 \leqslant s \leqslant n}$ of pairwise disjoint graphs isomorphic to α_s-components of G, $K_s, v_s^-, v_s^+ \simeq G_{\alpha_s}, g_s^-, g_s^+$ for $1 \leqslant s \leqslant n$.

This sequence is called *simple* if, for all $1 < s < n$, the connected components in K_s of v_s^- w.r.t. α_{s-1} and of v_s^+ w.r.t. α_{s+1} are disjoint.

In terms of the isomorphic representation of K_s, v_s^-, v_s^+ as $G_{\alpha_s}, g_s^-, g_s^+$, simplicity means that the α_{s-1}-component of g_s^- is disjoint from the α_{s+1}-component of g_s^+ in G_{α_s}, or that $(g_s^-)^{-1} \circ g_s^+ \notin G_{\alpha_{s-1}\cap\alpha_s} \circ G_{\alpha_s\cap\alpha_{s+1}}$. It implies that the *merged chain* obtained as

$$\sum_{s=1}^{n} (K_s, v_s^-, v_s^+) := K_1 \;\frac{v_1^+ = v_2^-}{\alpha_1\alpha_2}\; K_2 \quad \cdots \quad \frac{v_{n-1}^+ = v_n^-}{\alpha_{n-1}\alpha_n}\; K_n$$

is an E-graph. The simplicity condition also rules out inclusion relationships between the colour classes of K_s and K_{s+1} (other than at the ends, where an inclusion results in a trivial absorption). If $K_s \simeq G_{\alpha_s}$ then $\alpha_{s+1} \subseteq \alpha_s$ rules out a continuation beyond K_{s+1}, and the merger between K_s and K_{s+1} is trivial in the sense that it is isomorphic to just K_s.

The merged chains of simple sequences to be considered in the following will typically be of the form that $\alpha_s = \alpha \cap \beta_s$ for some sequence of subsets $\beta_s \subseteq E$ and a fixed subset $\alpha \subseteq E$. For simplicity we shall often write just $\alpha\beta$ instead of $\alpha \cap \beta$, especially when speaking of components and subgroups w.r.t. $\alpha \cap \beta$. E.g., $G_{\alpha\beta}$ stands for $G_{\alpha\cap\beta}$.

Definition 4. Let $G' \subseteq G$ be any subgroup, $\alpha \subseteq E$. We say that G' *admits chains of $(G_{\alpha\beta})_{\beta\subseteq E}$ up to length N*, if K is compatible with G' for every graph K obtained as the merged chain of a simple sequence of length up to N of components of the form $G_{\alpha\beta}$ for $\beta \subseteq E$.

2.2. Avoiding not just short cycles

Discounted lengths of cycles We want to measure the length of certain cycles in E-graphs in such a way as to reflect distances that discount repeated moves within the same $\alpha \subseteq E$. We present these notions in terms of Cayley groups but they could analogously be introduced in terms of E-graphs. We deal with cyclic words w of group elements, i.e., words $w = g_0 \cdots g_{n-1} = (g_t)_{t \in \mathbb{Z}_n}$, cyclically indexed modulo n.

Definition 5. Let G be a group with generator set E, with subgroups G_α for subsets $\alpha \subseteq E$ as above. A *non-trivial coloured cycle of length* n in G is any cyclic tuple $(g_t)_{t \in \mathbb{Z}_n}$ in G together with a colouring $\sigma \colon \mathbb{Z}_n \to \mathcal{P}(E)$ such that

(i) $\prod_{t \in \mathbb{Z}_n} g_t = g_0 \circ \cdots \circ g_{n-1} = 1$,

(ii) $g_t \in G_{\sigma(t)}$,

(iii) $g_t \notin G_{\sigma(t-1)\sigma(t)} \circ G_{\sigma(t)\sigma(t+1)}$.

G is called N-*acyclic* if all subgroups G_α for $\alpha \subseteq E$ are compatible with G and G has no non-trivial coloured cycles of lengths $n \leqslant N$.

The point of this notion is the way in which lengths of cycles in the Cayley graph of G are measured: we effectively count factors in subgroups G_α rather than the length of generator sequences that produce these factors. Therefore, the usual graph theoretic length of a coloured cycle of length n is a priori unbounded in terms of the underlying cycle of generator edges.

Condition (iii) concerns a property of the factors g_t in the subgroups $G_{\sigma(t)}$: it says that within this subgroup g_t is not equal to any product of two elements from the two subgroups $G_{\sigma(t)\sigma(t\pm 1)} \subseteq G_{\sigma(t)}$. Intuitively, this condition says that the effect of factor g_t cannot be absorbed via variations in the immediate predecessor and successor factors; this closely matches the condition on simple chains in Definition 3. Together with (ii), (iii) also rules out inclusions between adjacent colour classes: $\sigma(t) \not\subseteq \sigma(t \pm 1)$.

Lemma 6. *Let G be a group with generator set E, $k \in \mathbb{N}$. For every $\alpha \subseteq E$ with $|\alpha| < k$, let the subgroup G_α*

(a) admit chains of $(G_{\alpha\beta})_{\beta \subseteq E}$ up to length N, and

(b) have no non-trivial coloured cycles of length up to N.

Then there is a finite group G^* with generator set E s.t.:

(i) for every $\alpha \subseteq E$ with $|\alpha| < k$, $G^*_\alpha \simeq G_\alpha$,

and for all $\alpha \subseteq E$ with $|\alpha| \leqslant k$, the subgroups G^*_α

(ii) admit chains of $(G_{\alpha\beta})_{\beta \subseteq E}$ up to length N, and by (i) therefore also chains of $(G^*_{\alpha\beta})_{\beta \subseteq E}$, and

(iii) have no non-trivial coloured cycles of length up to N.

Compare Definition 1 and Remark 2 for the following.

Remark 7. *In the special case that $k = |E|$ and for $\alpha = E$, (ii) implies in particular that G^* is compatible with its β-components for all $\beta \subseteq E$. Because $G^* \simeq \mathrm{sym}(G^*)$, it follows that G^* is compatible with its subgroups G^*_α for $\alpha \subseteq E$ and reflects intersections.*

Proof of the lemma. We construct G^* as $G^* := \mathrm{sym}(H)$ for a graph $H = G \cup H^0$ consisting of the disjoint union of the Cayley graph of G and certain merged chains of components of G.

Consider any simple sequence $(K_s, v_s^-, v_s^+)_{1 \leqslant s \leqslant n}$ of length $n \leqslant N$ of components $K_s, v_s^-, v_s^+ \simeq G_{\alpha\beta_s}, g_s^-, g_s^+$ with $|\alpha| \leqslant k$. For any such sequence, we put the corresponding merged chain

$$\sum_{s=1}^{n} (K_s, v_s^-, v_s^+) := K_1 \;\frac{v_1^+ = v_2^-}{\alpha\beta_1\beta_2}\; \cdots \;\frac{v_{n-1}^+ = v_n^-}{\alpha\beta_{n-1}\beta_n}\; K_n \qquad (*)$$

as a separate connected component in H^0.

By construction, $G^* = \mathrm{sym}(G \cup H^0)$ admits chains of $(G_{\alpha\beta})_{\beta \subseteq E}$ up to length N as required (condition (ii), first formulation). Together with (i) this implies that G^* admits chains of $(G^*_{\alpha\beta})_{\beta \subseteq E}$ (condition (ii), second formulation) for the following reason. If the chain in question is such that all components $G^*_{\alpha\beta}$ have $|\alpha \cap \beta| < k$, (i) tells us that $G^*_{\alpha\beta} \simeq G_{\alpha\beta}$. If on the other hand some component $G^*_{\alpha\beta}$ has $|\alpha \cap \beta| = k$, then it must be that $|\alpha| = k$ and $\beta \supseteq \alpha$ and the merged chain is isomorphic to G^*_α; so in this case the claim boils down to G^*_α admits G^*_α, which is trivially true.

Towards (i) we claim that each one of the new connected components K as in $(*)$ is compatible with all $G_{\alpha'}$ for $|\alpha'| < k$. Let K as in $(*)$ and fix some $|\alpha'| < k$. Compatibility of K with $G_{\alpha'}$ depends only on the isomorphism types of α'-components of K. Every such component is obtained as a merged chain of a simple sequence of components of type $G_{\alpha'\alpha\beta_s}$ for s from some sub-interval of $[1, n]$. Since $|\alpha'| < k$, assumption (a) implies that this component is compatible with $G_{\alpha'}$.

It follows that $G^* = \mathrm{sym}(G \cup H^0)$ is compatible with all $G_{\alpha'}$ for $|\alpha'| < k$, whence $G^*_{\alpha'} \simeq G_{\alpha'}$ for $|\alpha'| < k$.

For (iii) it remains to argue that G^*_α does not have non-trivial coloured cycles of lengths $n \leqslant N$ whenever $|\alpha| \leqslant k$. Let $|\alpha| \leqslant k$ and let $((h_t)_{t \in \mathbb{Z}_n}, \sigma)$ be a non-trivial coloured cycle in G^*_α. We need to show that $n > N$.

Since $\sigma(t) \not\supseteq \sigma(t-1)$ (as a consequence of condition (iii) of Definition 5), $\sigma(t) \subsetneq \alpha$, whence $|\sigma(t)| < k$, for all t. Let $h_t = [u_t]^{G^*_\alpha}$ for a word u_t over $\sigma(t) \subseteq \alpha$, and put $w := u_1 \cdots u_n$. We want to show that $\prod_t h_t = [w]^{G^*_\alpha} \neq 1$ if $n \leqslant N$. It suffices to find an element of H on which w does not act as the identity. An element in a component of H^0 obtained as a suitable merged chain of components $G_{\alpha\sigma(t)} = G_{\sigma(t)}$ will serve this purpose. We look at the sequence

$$K_s, v_s^-, v_s^+ \simeq G_{\alpha\sigma(s)}, g_s^-, g_s^+ \simeq G^*_{\alpha\sigma(s)}, g_s^-, g_s^+$$

with $g_s^- := 1$ and $g_s^+ := [u_s]^G$ for $s \in Z_n$.

The sequence of these K_s, v_s^-, v_s^+ is simple in the sense of Definition 3, by condition (iii) in Definition 5. Therefore the corresponding merged chain $K := \sum_s (K_s, v_s^-, v_s^+)$ is a component of H provided $n \leqslant N$. But the element corresponding to $1 \in K_1$ is mapped by $[w]^{G^*}$ to the element corresponding to $g_n^+ \in K_n$, which is distinct from all elements represented in the components K_s for $s < n$ and in particular from $1 \in K_1$. It follows that, if $n \leqslant N$, $[w]^{G^*} \neq 1$, so that $(h_t)_{t \in \mathbb{Z}_n}$ cannot be a cycle in G^*_α. □

By iterated application of the lemma starting with $k = 1$ such that conditions (a) and (b) are trivially fulfilled (for $\alpha = \emptyset$!), we obtain the following, which technically is one of our key results. Compare Definition 5 for the following.

Corollary 8. *For every finite set E and every $N \in \mathbb{N}$ there is a finite N-acyclic group with generator set E.*

3. Hypergraph covers of bounded acyclicity

Basic notions A hypergraph is a structure $\mathfrak{A} = (A, S)$ consisting of a (finite) universe A together with a set of hyperedges $S \subseteq \mathcal{P}(A)$. The *width* of $\mathfrak{A}$ is the maximal cardinality among its hyperedges. With the hypergraph $\mathfrak{A} = (A, S)$ we associate its *Gaifman graph*, which is an undirected graph over the vertex set A with edges linking any pair of distinct vertices that are members of the same hyperedge $s \in S$ (a clique for every hyperedge of $\mathfrak{A}$). The notion of (induced) sub-hypergraph is the natural one: think of removing all elements not in the designated subset from both the universe and from every hyperedge.

Acyclicity in hypergraphs A hypergraph is *conformal* if every clique in its Gaifman graph is contained in some hyperedge; in analogy with guardedness in relational structures, cf. Section 4, we also say that every clique must be *guarded* by a hyperedge. More generally, a set of nodes is guarded if it is contained in some hyperedge.

A hypergraph (or its Gaifman graph) is *chordal* if every cycle of length greater than 3 in the Gaifman graph has a chord. We use cyclic words $(a_t)_{t \in \mathbb{Z}_n}$ to denote cycles (indexing modulo n); this cycle is *chordless* if $\{a_i, a_j\}$ is not guarded unless $i = j, j \pm 1$.

It is known from classical hypergraph theory, cf. [5, 4], that a hypergraph is tree-decomposable (also called acyclic) if, and only if, it is both conformal and chordal. $\mathfrak{A} = (A, S)$ is tree-decomposable if it admits a *tree decomposition* $\mathcal{T} = (T, \delta)$: T is a tree and $\delta\colon v \mapsto \delta(v) \in S$ maps the nodes of T to hyperedges of $\mathfrak{A}$ in such a manner that $\mathrm{im}(\delta) = S$ and, for every node $a \in \mathfrak{A}$, the subset $\{v \in T\colon a \in \delta(v)\}$ is connected in T. An equivalent characterisation requires that $\mathfrak{A}$ can be reduced to the empty hypergraph by repeated application of two kinds of reduction steps: removal of a node that is covered by at most one hyperedge, and removal of a hyperedge that is fully contained in some other hyperedge.

The bounded variants of acyclicity and its constituents, which are relevant to us, are the following. The N-bounded version of each one of these properties precisely captures the requirement that every induced sub-hypergraph of size up to N has the unqualified property.

Definition 9. Let $N \in \mathbb{N}$. A hypergraph $\mathfrak{A}$ is called *N-conformal* (*N-chordal*) if it does not have any unguarded cliques up to size N (chordless cycles up to length N). It is called *N-acyclic* if it is N-conformal and N-chordal.

Hypergraph covers The notion of hypergraph bisimulation is the natural generalisation of bisimulation between graph-like structures (transition systems). It captures the idea of a back-and-forth correspondence whose individual matches are bijections between individual hyperedges and whose back-and-forth requirements ensure that the overlap patterns between hyperedges in one hypergraph can be simulated in the other. In this sense hypergraph bisimulation is also at the combinatorial core of guarded bisimulation (as if stripped of the relational information within relational hyperedges). Here we discuss a special case, viz. hypergraph bisimulations induced by a hypergraph homomorphism from one (covering) hypergraph onto another.

Definition 10. A map $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ between hypergraphs $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ and $\mathfrak{A} = (A, S)$ is a *hypergraph homomorphism* if $\pi \restriction \hat{s}$ is a bijection between the hyperedge $\hat{s}$ and its image $\pi(\hat{s}) \in S$, for every $\hat{s} \in \hat{S}$.

A homomorphism $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is a (bisimilar) *hypergraph cover* if it satisfies the following *back*-property: for every $s \in S$ there is some $\hat{s} \in \hat{S}$ such that $\pi(\hat{s}) = s$ and, whenever $\pi(\hat{s}) = s$ and $s' \in S$, then there is some $\hat{s}' \in \hat{S}$ such that $\pi(\hat{s}') = s'$ and $\pi(\hat{s} \cap \hat{s}') = s \cap s'$.

Note that the homomorphism requirement for covers settles the *forth*-property in the back-and-forth view.

The (conformal and) N-acyclic hypergraph covers to be constructed below are hypergraph covers $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ by (conformal and) N-acyclic hypergraphs $\hat{A}$.

3.1. Millefeuilles of hypergraphs

Let $\mathfrak{A} = (A, S)$ be a finite hypergraph, and let the colours $e \in E$ be associated with guarded subsets of $\mathfrak{A}$ through a map $\rho\colon e \mapsto \rho(e) \subseteq A$. We consider stacks of copies of the hypergraph $\mathfrak{A}$ that are selectively joined in the subsets $\rho(e)$. For $a \in A$, let $\alpha_a := \{e \in E\colon a \in \rho(e)\}$.

For a group G with generator set E, we write G_a for the subgroup generated by α_a; $G_{aa'}$ for the subgroup generated by $\alpha_{aa'} := \alpha_a \cap \alpha_{a'} = \{e \in E\colon a, a' \in \rho(e)\}$, etc.

On $A \times G$ consider the equivalence relation

$$(a, g) \approx (a, g') \quad :\Leftrightarrow \quad g^{-1} \circ g' \in G_a.$$

We write $[a, g]$ for the equivalence class of (a, g) w.r.t. $\approx$, and lift this notation to tuples and sets of elements as, e.g., in $[s, g] := \{[a, g] \colon a \in s\}$. We put

$$\mathfrak{A} \times_E G := (\hat{A}, \hat{S}) \quad \text{with} \quad \begin{cases} \hat{A} := (A \times G)/\approx, \\ \hat{S} := \{[s, g] \colon s \in S, g \in G\}. \end{cases}$$

The definitions of $\approx$ and $\hat{S}$ imply that $[a, g] \in [s, h]$ iff $a \in s$ and $g^{-1} \circ h \in G_a$ iff $a \in s$ and $[a, g] = [a, h]$.

Note that $\approx$ is trivial in restriction to $A \times \{g\}$, whence $(A \times \{g\})/\approx$ is naturally identified with $A \times \{g\}$ and carries the hypergraph structure of $\mathfrak{A}$. We refer to the isomorphic copies of $\mathfrak{A}$ thus embedded as induced hypergraphs as to the *layers* of $\mathfrak{A} \times_E G$, denoted $\mathfrak{A} \times \{g\}$. The natural projection $\pi \colon \mathfrak{A} \times_E G \to \mathfrak{A}$ is a cover.

Proposition 11. *Let G be N-acyclic with generator set E, $\mathfrak{A} \times_E G$ as above.*

(i) Any chordless cycle of length up to N in $\mathfrak{A} \times_E G$ must be contained within a single layer of $\mathfrak{A} \times_E G$.

(ii) Any unguarded clique of size up to N in $\mathfrak{A} \times_E G$ must be contained within a single layer of $\mathfrak{A} \times_E G$.

So N-conformality and N-chordality are preserved in the passage to $\mathfrak{A} \times_E G$.

If $\mathfrak{A}$ is conformal and of width $w \leqslant N$, then $\mathfrak{A} \times_E G$ is also conformal.

Proof. For the proof of (i) assume that $(\hat{a}_t)_{t \in \mathbb{Z}_n}$ is a chordless cycle of length $3 < n \leqslant N$ in $\hat{\mathfrak{A}} = \mathfrak{A} \times_E G = (\hat{A}, \hat{S})$. We let $\hat{s}_t = [s_t, h_t] \in \hat{S}$ be a sequence of linking hyperedges such that $\hat{a}_t \in \hat{s}_t \cap \hat{s}_{t+1}$ and assume that the (s_t, h_t) are chosen such that the number of jumps between distinct layers $\mathfrak{A} \times \{h_t\}$ is minimal: with $J := \{t \colon h_t \neq h_{t+1}\}$, the (s_t, h_t) have been chosen so as to minimise $|J|$. Clearly $0 \leqslant |J| \leqslant n$, and our goal is to show that this minimisation implies $J = \emptyset$. Put

$$u_t := h_t^{-1} \circ h_{t+1}$$

and let $\sigma(t) := \alpha_{a_t} = \{e \in E \colon a_t \in \rho(e)\}$. Then $\hat{a}_t \in \hat{s}_t \cap \hat{s}_{t+1}$ implies that $u_t \in G_{\sigma(t)}$. Clearly $u_t \neq 1$ iff $t \in J$, and $\prod_t u_t = \prod_{t \in J} u_t = 1$.

That the cycle of the $\hat{a}_t$ is chordless implies that, for $t' \neq t \pm 1$, $\hat{a}_t$ and $\hat{a}_{t'}$ cannot be represented in the same layer of $\hat{\mathfrak{A}}$ or that $(a_t, a_{t'})$ are not members of a common hyperedge of $\mathfrak{A}$; because the sets $\rho(e)$ are guarded, the latter case implies that $\sigma(t) \cap \sigma(t') = \emptyset$ and thus $G_{\sigma(t)\sigma(t')} = \{1\}$.

We claim that, for non-empty J, $(u_t)_{t \in J}$ would be a non-trivial coloured cycle in G, coloured by the natural restriction of σ. For this we verify that any violation of condition (iii) in Definition 5 would allow us to eliminate one of the remaining jumps, contradicting the minimality of $|J|$.

Consider next neighbours $t' < t$ in J along this cycle. Since there are no jumps between t' and t, $\hat{a}_{t'}$ and $\hat{a}_t$ are both represented in layer $\mathfrak{A} \times \{h_t\}$, which implies $\sigma(t) \cap \sigma(t') = \emptyset$ and $G_{\sigma(t)\sigma(t')} = \{1\}$ unless $t = t' + 1$.

Assume then that $t' < t < t''$ are next neighbours in J and that – contrary to condition (iii) in Definition 5 – we had $u_t \in G_{\sigma(t')\sigma(t)} \circ G_{\sigma(t)\sigma(t'')}$.

Clearly this implies that $t' = t - 1$ or $t'' = t + 1$, since otherwise $G_{\sigma(t')\sigma(t)} = G_{\sigma(t)\sigma(t'')} = \{1\}$ while $u_t \neq 1$.

Suppose first that, e.g., $t' = t - 1$ but $t'' \neq t + 1$. Then $u_t \in G_{\sigma(t')\sigma(t)} \circ G_{\sigma(t)\sigma(t'')} = G_{\sigma(t')\sigma(t)} \subseteq G_{\sigma(t')}$ implies that $\hat{a}_{t'}$ is also represented in layer $\mathfrak{A} \times \{h_{t''}\}$, contradicting minimality of $|J|$.

If $t' = t - 1$ and $t'' = t + 1$, then we may use a decomposition of $u_t \in G_{\sigma(t')\sigma(t)} \circ G_{\sigma(t)\sigma(t'')}$ as $u_t = g' \circ g''$ with $g' \in G_{\sigma(t')\sigma(t)}$ and $g'' \in G_{\sigma(t)\sigma(t'')}$ to find representations of $\hat{a}_{t'}, \hat{a}_t, \hat{a}_{t''}$ in the common layer $\mathfrak{A} \times \{g\}$ for $g = h_t \circ g' = h_{t''} \circ (g'')^{-1}$.

The proof of (ii) is similar in spirit, looking at a clique $(\hat{a}_t)_{t \in \mathbb{Z}_n}$ of minimal size that is not contained in a single layer.

For the last claim of the proposition, note that cliques in $\mathfrak{A} \times_E G$ project injectively onto cliques in $\mathfrak{A}$; if $\mathfrak{A}$ is conformal, then $\mathfrak{A}$ and $\mathfrak{A} \times_E G$ cannot have cliques larger than w. $\square$

Local covers In a preparatory step we want to obtain L-local N-acyclic covers, which will then be stacked and glued by means of the construction indicated above. We write $N^\ell(a)$ for the Gaifman neighbourhood of radius ℓ of a, consisting of all nodes at distance up to ℓ from a in the Gaifman graph.

Definition 12. Let $L \in \mathbb{N}$. A homomorphism $\pi \colon \mathfrak{B} \to \mathfrak{A}$ between hypergraphs is called an *L-local cover* at $a \in \mathfrak{A}$ if for some $b \in \pi^{-1}(a)$, π satisfies the *back*-condition for bisimilar covers as far as extensions at hyperedges in $N^{L-1}(b)$ are concerned: if $\hat{s} \subseteq N^{L-1}(b)$ and $s = \pi(\hat{s})$ and $s' \in S$ are such that $s \cap s' \neq \emptyset$, then there is some $\hat{s}' \in \hat{S}$ such that $\pi(\hat{s}') = s'$ and $\pi(\hat{s} \cap \hat{s}') = s \cap s'$.

The construction of these will rely on the availability of (full rather than local) N-acyclic and conformal covers of hypergraphs of smaller width. The basic step in the construction is reflected in the following simple observations. For technical reasons we assume that the set of hyperedges is closed under subsets.

Consider a node a in a hypergraph $\mathfrak{A} = (A, S)$. The *localisation* of $\mathfrak{A}$ at a is the hypergraph $\mathfrak{A} \upharpoonright N^1_*(a)$ induced by S on the subset $N^1_*(a) := N^1(a) \setminus \{a\}$. Its hyperedges are the intersections of hyperedges $s \in S$ with $N^1_*(a)$. Note that for conformal $\mathfrak{A}$, every $s \cap N^1(a)$ is contained in some hyperedge s' with $a \in s'$. For conformal $\mathfrak{A}$, therefore, the width of $\mathfrak{A} \upharpoonright N^1_*(a)$ is strictly less than that of $\mathfrak{A}$.

Proofs of the observation and the following two lemmas are omitted.

Observation 13. *Let $a \in \mathfrak{A} = (A, S)$ be conformal and consider a cover $\pi\colon \mathfrak{B}_0 \to \mathfrak{A} \restriction N^1_*(a)$ with $\mathfrak{B}_0 = (B_0, T_0)$. For a new element $b \notin B_0$, let $B := B_0 \cup \{b\}$ and extend π by $\pi(b) := a$. Then the hypergraph $\mathfrak{B} := (B, T)$ with $T := \{t \subseteq B : t \setminus \{b\} \in T_0, \pi(t) \in S\}$ provides a cover of $\mathfrak{A} \restriction N^1(a)$. Moreover,*

(i) if $\mathfrak{B}_0$ is (N-)conformal, then so is $\mathfrak{B}$.

(ii) if $\mathfrak{B}_0$ is N-chordal, then so is $\mathfrak{B}$.

In order to enlarge the radius of local covers based on this idea, we first discuss a simple glueing mechanism that preserves acyclicity and conformality.

Lemma 14. *Let $\pi_0\colon \mathfrak{B}_0 \to \mathfrak{A}$ a homomorphism that bijectively maps hyperedges of $\mathfrak{B}_0$ onto hyperedges of $\mathfrak{A}$, and let $\rho\colon \mathfrak{C} \to \mathfrak{A}$ be a cover. Then there is a cover $\pi\colon \mathfrak{B} \to \mathfrak{A}$ extending π_0 in the sense that $\mathfrak{B} \supseteq \mathfrak{B}_0$ and $\pi_0 = \pi \restriction B_0$. Moreover:*

(i) if $\mathfrak{B}_0$ and $\mathfrak{C}$ are (N-)conformal, then so is $\mathfrak{B}$.

(ii) if $\mathfrak{B}_0$ and $\mathfrak{C}$ are N-chordal, then so is $\mathfrak{B}$.

Lemma 15. *Suppose that N-acyclic, conformal covers are available for all width k hypergraphs. Then there is, for every hypergraph $\mathfrak{A}$ of width $k+1$, every element $a \in A$ and every $L \in \mathbb{N}$, an L-local cover $\pi\colon \mathfrak{B}, b \to \mathfrak{A}, a$ at a by an N-acyclic and conformal hypergraph $\mathfrak{B}$.*

Availability of N-acyclic covers for width 2 hypergraphs follows from [10]. Note that width 2 hypergraphs are graphs $\mathfrak{A} = (A, E)$, and the basic construction of Cayley groups from graphs (in this case, from regularly E-coloured trees of depth N) as indicated in Section 2.1 can be used to obtain a Cayley group G of girth greater than N whose set of involutive generators is the set E of edges of $\mathfrak{A}$. Then the product $\mathfrak{A} \otimes G$ with vertices $(a, g) \in A \times G$ and edges of the form $\{(a, g), (a', g \circ e)\}$ above edge $e = \{a, a'\} \in E$ provides an N-acyclic cover $\pi\colon \mathfrak{A} \otimes G \to \mathfrak{A}$. This settles the base case for the inductive application of the lemma to the construction of N-acyclic covers of finite hypergraphs of any width.

From local to global covers Suppose $\mathfrak{A} = (A, S)$ and $a_0 \in \mathfrak{A}$ and $S_0, S_1 \subseteq S$ are such that

$$\bigcup S_0 \subseteq A \setminus \bigcup S_1 \subseteq N^{L-1}(a_0)$$

and $d(\bigcup S_0, \bigcup S_1) > N$.

Think of S_0 as the core region of some L-local cover of a given hypergraph that is such that every hyperedge of that original hypergraph is covered by some $s \in S_0$; the set S_1, on the other hand, comprises all those hyperedges in the periphery of this local cover, which may still be lacking responses to *back*-requirements. Missing hyperedge neighbours of peripheral hyperedges are to be supplied through glueing with hyperedges in the core region of new copies of $\mathfrak{A}$. For this we need a surplus of core hyperedges compared to the demands created by the peripheral hyperedges. It is to this end that stacking is used: to create many layers of copies of core hyperedges without unduly increasing the number of peripheral ones.

In the given situation, the glueing of isomorphic copies of $\mathfrak{A}$ is achieved with $\hat{\mathfrak{A}} = \mathfrak{A} \times_E G$, where $E = \{1, \ldots, K\} \times S_0$ and $\rho\colon (i, s) \mapsto s \subseteq A$. As before, we let G be an N-acyclic group with generator set E.

$\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ is a cover of $\mathfrak{A}$ w.r.t. the natural projection $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$, such that all the copies of $s \in S_0$ in the different layers of $\hat{\mathfrak{A}}$ are far from each other and far from the copies of elements $s \in S_1$. Moreover, the multiplicity ratio between centre and boundary is improved at least by a factor of K. On one hand, $\hat{\mathfrak{A}}$ has $|G|$ many disjoint isomorphic copies of $\mathfrak{A} \restriction \bigcup S_0 \subseteq \mathfrak{A} \restriction (A \setminus \bigcup S_1)$, because these regions are far from any glueing sites. For $s \in S_1$, on the other hand, the number of distinct covers $[s, g]$ above s is at most $|G| : |G_s|$, where G_s is the subgroup generated by $\{(i, s) : 1 \leqslant i \leqslant K\}$ and therefore has at least K elements. This is because $[s, g] = [s, g']$ whenever $g^{-1} \circ g \in G_s$.

Choosing $K > |S_1|$, there is an injection κ from hyperedges $\hat{s}$ of $\hat{\mathfrak{A}}$ above S_1 into layers of $\hat{\mathfrak{A}}$: the number of such hyperedges $\hat{s}$ is bounded by $|S_1||G|/K < |G|$.

Let now $\pi_0\colon \mathfrak{A} \to \mathfrak{A}_0$ be an L-local cover of $\mathfrak{A}$ at $\pi(a_0)$ by some conformal and N-acyclic $\mathfrak{A}$; let $S_0, S_1 \subseteq S$ be as above and such that for every $s \in S_1$ there is some $s' \in S_0$ such that $\pi_0(s) = \pi_0(s')$ – we fix such a selection of s' for every $s \in S_1$. Let further $\pi\colon \hat{\mathfrak{A}} = \mathfrak{A} \times_E G \to \mathfrak{A}$ be constructed for $S_0, S_1 \subseteq S$ as above, with $K > |S_1|$ and an injection κ from $\pi^{-1}(S_1)$ into G. Clearly $\hat{\pi}\colon \hat{\mathfrak{A}} \to \mathfrak{A}_0$, $\hat{\pi} := \pi_0 \circ \pi$, is an L-local cover by a conformal and N-acyclic hypergraph. We may then construct a full conformal and N-acyclic cover $\tilde{\pi}\colon \tilde{\mathfrak{A}} \to \mathfrak{A}_0$ as follows.

The hypergraph $\tilde{\mathfrak{A}}$ is obtained from $\hat{\mathfrak{A}}$ simply by identifying $\hat{s} \in \pi^{-1}(S_1)$ with $[s', \kappa(\hat{s})] \in \hat{S}_0$. As this identification is compatible with $\hat{\pi}$, we can choose $\tilde{\pi}$ to be the natural projection induced by $\hat{\pi}$. It is obvious that $\tilde{\pi}\colon \tilde{\mathfrak{A}} \to \mathfrak{A}_0$ is a full cover since all defects in $\pi_0\colon \mathfrak{A} \to \mathfrak{A}_0$ have been healed through the glueing of peripheral with central hyperedges. It also not hard to see that the identifications between $\hat{s} \in \pi^{-1}(S_1)$ with $[s', \kappa(\hat{s})] \in \hat{S}_0$ do not violate conformality or N-chordality. We have thus obtained our second main technical result.

Theorem 16. *For every $N \in \mathbb{N}$, every finite hypergraph admits a bisimilar cover by some finite conformal and N-acyclic hypergraph.*

4. Two applications to the guarded fragment

We deal with the guarded fragment GF $\subseteq$ FO as introduced in [2]; we assume some familiarity with its role as a versatile analogue of modal logic in the much richer setting of arbitrary relational structures. For background and key results in its model theory see in particular [7].

The key feature of GF is its relativised quantification pattern, which only allows quantification over *guarded tuples*, i.e., tuples that are covered by some ground atom. Instead of arbitrary FO quantification, GF admits only *guarded quantification* of the form

$$\exists \mathbf{y}.\alpha(\mathbf{x})\varphi(\mathbf{x}) \quad \text{and, dually} \quad \forall \mathbf{y}.\alpha(\mathbf{x})\varphi(\mathbf{x})$$

where $\mathbf{y} \subseteq \mathbf{x}$ is a tuple of variables among those that occur in the *guard* $\alpha(\mathbf{x})$, which is an atom in which all the free variables of φ *must* occur. Here we use the shorthand $\exists \mathbf{y}.\alpha\varphi$ for $\exists \mathbf{y}(\alpha \wedge \varphi)$ and $\forall \mathbf{y}.\alpha\varphi$ for $\forall \mathbf{y}(\alpha \rightarrow \varphi)$.

With a relational τ-structure $\mathfrak{A}$ with universe A and relations $(R^{\mathfrak{A}})_{R\in\tau}$ we associate the *hypergraph of guarded sets* $(A, S[\mathfrak{A}])$, whose hyperedges are precisely all singleton sets together with all the subsets of sets $\{a \colon a \in \mathbf{a}\}$ for $\mathbf{a} \in R^{\mathfrak{A}}, R \in \tau$. Clearly the width of this induced hypergraph is bounded by the width of the signature τ (the maximal arity in τ).

Also note that the Gaifman graph of $\mathfrak{A}$ *is* the Gaifman graph associated with the hypergraph of guarded sets. Through $(A, S[\mathfrak{A}])$, hypergraph theoretic notions like conformality or N-chordality transfer naturally to relational structures $\mathfrak{A}$.

The natural back-and-forth equivalence that captures the restricted nature of guarded quantification is *guarded bisimulation*, which we denote as $\sim_{\mathrm{g}}$. Its finite approximations $\sim^{\ell}_{\mathrm{g}}$ correspond to equivalence $\equiv^{\mathrm{GF}}_{\ell}$ w.r.t. GF up to nesting depth ℓ. GF is preserved under $\sim_{\mathrm{g}}$ and in fact expressively complete for $\sim_{\mathrm{g}}$-invariant FO. It is this characterisation theorem of GF – classically due to by Andréka, van Benthem and Németi [2] – that we aim to prove in the sense of finite model theory in Section 4.2 below.

We may think of guarded bisimulations between τ-structures $\mathfrak{A}$ and $\mathfrak{B}$ as hypergraph bisimulations between the hypergraphs $(A, S[\mathfrak{A}])$ and $(B, S[\mathfrak{B}])$, in which we additionally require the local bijections between hyperedges (guarded sets) to be local isomorphisms of the relational structures. Correspondingly, we define guarded covers $\pi \colon \hat{\mathfrak{A}} \rightarrow \mathfrak{A}$ of one τ-structure by another to be relational homomorphisms that are covers w.r.t. the associated hypergraphs of guarded sets.

Importantly, any hypergraph cover $\pi \colon (\hat{A}, \hat{S}) \rightarrow (A, S[\mathfrak{A}])$ induces a unique relational structure $\hat{\mathfrak{A}}$ on the universe $\hat{A}$ that turns π into a guarded cover. From this we obtain the following as a direct corollary to Theorem 16.

Corollary 17. *For every $N \in \mathbb{N}$, every finite relational structure admits a guarded bisimilar cover by some finite conformal and N-acyclic structure.*

The following generalises the finite model property of GF [7] and its strengthening in [3].

Corollary 18. GF *has the finite model property in restriction to any class of relational structures that is defined in terms of finitely many forbidden cyclic configurations.*

That this cannot be strengthened to arbitrary choices of finitely many forbidden configurations follows from the undecidability of GF with functionality constraints [7].

4.1. Structures of bounded acyclicity

In the following we sketch some basic analysis of hypergraphs and relational structures of bounded acyclicity. We shall often refer to *sufficiently acyclic* hypergraphs or structures to appeal to some not necessarily explicitly specified bound N such that corresponding constructions go through for all conformal and N-acyclic hypergraphs, or structures whose Gaifman graph is conformal and N-acyclic. Uniformity of a suitable bound N in explicitly specified parameters is always understood.

Besides the parameter N specifying the global acyclicity requirements for all $\mathfrak{A}$ under consideration, we often deal with a locality parameter n to say which Gaifman distances and path lengths are currently considered as *short*. In typical game arguments, for instance, n will be shrinking from round to round, with a dependency like $n_i = 2^{m-i}$ in round i of an m-round game. With a choice for n set, we often refer to *short* paths when we mean paths of length up to n.

A *chordless path* in $\mathfrak{A}$ is a chordless path in the Gaifman graph of $\mathfrak{A}$. Examples of (short) chordless paths are shortest paths between two nodes a and a' (at distance $d(a, a') \leqslant n$).

For the rest of this section we focus on the core concepts and methods without giving any proof details.

Short chordless paths and bounded convexity A first surprising feature in sufficiently acyclic $\mathfrak{A}$ is that the number of all nodes on shortest paths between two given nodes a and a' at distance $d(a, a') \leqslant n$ can be bounded in terms of the width w of $\mathfrak{A}$. Similarly, even the number of nodes on any short chordless paths between two nodes at short distance can be bounded – and this is strengthened even further to yield a corresponding bounded closure operator below.

Definition 19. (i) For $B \subseteq \mathfrak{A}$ we let $D(a, B)$ be the set of all nodes on shortest paths between a and B.

(ii) A subset $B \subseteq \mathfrak{A}$ is *n-closed* if any chordless path of length up to n between $a, a' \in B$ fully runs in B.

(iii) For $n \in \mathbb{N}$, the convex n-closure of a tuple $\mathbf{a}$ in $\mathfrak{A}$ is $\mathrm{cl}_n(\mathbf{a}) := \bigcap \{ B \subseteq \mathfrak{A} \colon \mathbf{a} \subseteq B \text{ } n\text{-closed} \}$.

It is not hard to see that, provided $\mathfrak{A}$ is sufficiently acyclic in relation to $d(a, B)$, the size of the set $D(a, B) \setminus B$ is bounded by the product of $d(a, B)$ and the width w of $\mathfrak{A}$: in fact $2n$-chordality implies that the subset D_k of elements at distance k from B in D must be a clique for $1 \leqslant k \leqslant d(a, B)$ if $d(a, B) \leqslant n$; hence, by conformality, each D_k is contained in a hyperedge and its size bounded by w.

It is considerably harder to show that also $\mathrm{cl}_n(B)$ is uniformly size bounded (in terms of $|B|$, w, n) in all sufficiently acyclic $\mathfrak{A}$. For this one establishes (by induction on w) the existence of *some* size-bounded n-closed superset.

Lemma 20. *For $n \in \mathbb{N}$ there is a function $f_n(w, k)$ such that, for all sufficiently acyclic $\mathfrak{A}$ of width w, every $\mathbf{a} \in A^k$ is contained in some n-closed subset $B(\mathbf{a})$ of size $\leqslant f_n(w, k)$. Hence $|\mathrm{cl}_n(\mathbf{a})| \leqslant f_n(w, k)$.*

The following is useful towards extension arguments.

Lemma 21. *Let $B \subseteq \mathfrak{A}$ be $2n$-closed and $a \in \mathfrak{A}$ such that $d(a, B) \leqslant n$. Let $\hat{B} := \mathrm{cl}_n(B \cup \{a\})$ and $D \subseteq B$ the subset of nodes directly linked to nodes in $\hat{B} \setminus B$. Then*

(i) $\hat{B} \setminus B$ is connected;

(ii) D is a clique.

It follows that $\mathrm{cl}_n(B \cup \{a\}) = \mathrm{cl}_n(B) \cup \mathrm{cl}_n(D \cup \{a\})$, as no edge within $\hat{B}$ can bridge D.

Free realisations of small convex configurations

Definition 22. Let $\mathfrak{A} = (A, S)$ be a hypergraph.

(i) For $s \in S$, $B \subseteq A$ and $t \subseteq s \cap B$, let $d_t(s, B)$ be the usual distance in $\mathfrak{A} \restriction (A \setminus t)$ between $s \setminus t$ and $B \setminus t$.

(ii) For $s \in S$, $B \subseteq A$ and $t \subseteq s \cap B$, we say that s and B are *n-free over t* if $d_t(s, B) > n$.

(iii) $\mathfrak{A}$ itself is called *(n, K)-free* if, for all $s \in S$ and $B \subseteq A$ of size $|B| \leqslant K$ and $t \subseteq s \cap B$, there is some $s' \in S$ such that s' and B are n-free over t and $s' \sim s$.

Lemma 23. *Let $n, K \in \mathbb{N}$ and let $\mathfrak{A} = (A, S)$ be a hypergraph. For sufficiently large M and N consider $E = \{0, \ldots, M\} \times \{t \subseteq s \colon s \in S\}$ with the association $\rho\colon (i, t) \mapsto t \subseteq A$ and let G be an N-acyclic group with generator set E. Then the hypergraph $\hat{\mathfrak{A}} := \mathfrak{A} \times_E G$ is an (n, K)-free cover of $\mathfrak{A}$.*

Let $B \subseteq \mathfrak{A}$ be connected and n-closed, i.e., such that $\mathrm{cl}_n(B) = B$, in a sufficiently free and acyclic structure $\mathfrak{A}$, where in particular $|B|$ is small enough to guarantee acyclicity of $\mathfrak{A} \restriction B$. Then $\mathfrak{A} \restriction B$ admits a tree decomposition by guarded subsets, $\mathcal{T} = (T, \delta)$ where $\delta\colon v \mapsto \delta(v) \in S[\mathfrak{A}]$. Let the guarded tuple $\mathbf{b} \in B$ be represented in the designated root λ of T, $\delta(\lambda) = \mathbf{b}$.

With $\mathcal{T}$ we associate a GF-formula $\varphi_{\mathcal{T}}(\mathbf{x}) := \varphi_{\mathcal{T},\lambda}$ describing the existential GF-type of $\big(\mathfrak{A} \restriction B, \mathbf{b}\big)$. Formulae $\varphi_{\mathcal{T},v}$ are defined by induction w.r.t. the depth of $v \in T$. For leaves v, $\varphi_{\mathcal{T},v}$ is a quantifier-free description of the atomic type of $\delta(v)$ in $\mathfrak{A}$. From formulae $\varphi_{\mathcal{T},v_i}(\mathbf{x}^{(i)})$ for the children v_i of $v \in T$ we obtain $\varphi_{\mathcal{T},v}$ in the obvious manner as a formula of the form $\chi(\mathbf{x}) \wedge \bigwedge_i \exists \mathbf{x}^{(i)}.\alpha^{(i)} \varphi_{\mathcal{T},v_i}(\mathbf{x}^{(i)})$ with guards $\alpha^{(i)}$ abstracted from v_i and a description $\chi(\mathbf{x})$ of the atomic type of $\delta(v)$ in $\mathfrak{A}$.

Lemma 24. *For $\mathfrak{A}$, $\mathcal{T}$, $\mathbf{b}$, $\varphi_{\mathcal{T}}(\mathbf{x})$ as above: if $\mathfrak{A}'$ is sufficiently acyclic and free, then $\mathfrak{A}', \mathbf{b}' \models \varphi_{\mathcal{T}}(\mathbf{x})$ implies that there is an n-closed subset $B' \subseteq \mathfrak{A}'$ such that*

$$\mathfrak{A} \restriction B, \mathbf{b} \simeq \mathfrak{A}' \restriction B', \mathbf{b}'.$$

Remark 25. *In the above situation, for any subtree $T_0 \subseteq T$ such that $B_0 := \bigcup\{\delta(v) \colon v \in T_0\}$ is n-closed, any isomorphism $\mathfrak{A} \restriction B_0 \simeq \mathfrak{A}' \restriction B_0'$ for n-closed $B_0' \subseteq \mathfrak{A}'$ extends to an isomorphism $\mathfrak{A} \restriction B \simeq \mathfrak{A}' \restriction B'$ with n-closed B'.*

4.2. The FMT characterisation theorem

We sketch a proof of the following finite model theory version of the classical characterisation of GF as the $\sim_{\mathrm{g}}$-invariant fragment of FO from [2].

Theorem 26. GF *precisely captures the guarded bisimulation invariant fragment of* FO *also in restriction to just finite relational structures:* FO/$\sim_{\mathrm{g}}$ $\equiv_{\mathrm{fin}}$ GF. *I.e., the following are equivalent for any sentence $\varphi \in$ FO(τ):*

(i) φ is preserved under $\sim_{\mathrm{g}}$ between finite τ-structures.

(ii) $\varphi \equiv_{\mathrm{fin}} \varphi'$ for some $\varphi' \in$ GF(τ).

Back-and-forth in free and acyclic models The crux in Theorem 26 is the proof of expressive completeness of GF for $\sim_{\mathrm{g}}$-invariant FO-sentences. This is achieved with an upgrading of suitable levels $\sim_{\mathrm{g}}^{\ell}$ of guarded bisimulation equivalence to levels $\equiv_q$ of elementary equivalence in suitably prepared models, which then shows that any φ as in the theorem is actually preserved under some $\sim_{\mathrm{g}}^{\ell}$. Then the usual Ehrenfeucht–Fraïssé techniques imply that φ is equivalent to the disjunction of those GF-sentences that characterise the full $\sim_{\mathrm{g}}^{\ell}$-types of models of φ; this yields the desired $\varphi' \in$ GF(τ) since there are only finitely many $\sim_{\mathrm{g}}^{\ell}$-types.

For the desired upgrading we provide an extension lemma, which will cover the crucial back-and-forth requirements of a single round in the first-order Ehrenfeucht–Fraïssé game. We rely on sufficient levels of acyclicity and freeness to make sure we can maintain the appropriate closure conditions.

For a partial bijection ρ between two τ-structures $\mathfrak{A}$ and $\mathfrak{A}'$ with domain $B \subseteq \mathfrak{A}$ and image $B' \subseteq \mathfrak{A}'$ we use the notation $\rho\colon B \longmapsto_{\mathrm{g}}^{q} B'$ to indicate that ρ is compatible with GF^q-equivalence or $\sim_{\mathrm{g}}^{q}$, in the sense that $\mathfrak{A}, \mathbf{b} \sim_{\mathrm{g}}^{q} \mathfrak{A}', \mathbf{b}'$ for all (guarded) tuples $\mathbf{b} \in B$ and $\mathbf{b}' = \rho(\mathbf{b}) \in B'$.

Lemma 27. *Let $Q \geqslant q + f_n(w, w+1)$, where f_n is the bound on sizes of n-closures from Lemma 20 and w the width of τ. Let $\mathfrak{A}$ and $\mathfrak{A}'$ be sufficiently free and acyclic,*

$$\rho \colon B \longmapsto^{Q}_{\mathrm{g}} B'$$

a bijection between $2n$-closed subsets $B = \mathrm{dom}(\rho) \subseteq \mathfrak{A}$ and $B' = \mathrm{im}(\rho) \subseteq \mathfrak{A}'$. Then there is, for every $a \in \mathfrak{A}$, an extension $\hat{\rho} \supseteq \rho$,

$$\hat{\rho} \colon \hat{B} \longmapsto^{q}_{\mathrm{g}} \hat{B}'$$

with $a \in \mathrm{dom}(\hat{\rho})$ and such that $\hat{B} = \mathrm{dom}(\hat{\rho})$ and $\hat{B}' = \mathrm{im}(\hat{\rho})$ are n-closed.

Proof. Let us work in the expansions $\mathfrak{A}_q$ and $\mathfrak{A}'_q$ of $\mathfrak{A}$ and $\mathfrak{A}'$ by predicates that mark the GF^q-types of guarded tuples. [In effect this means that ρ is compatible with GF^{Q-q} over the expansions, and we need $\hat{\rho}$ to be just a local isomorphism w.r.t. these expansions.]

If $d(a, B) > n$, pick any $a' \in \mathfrak{A}'$ such that $\mathfrak{A}_q, a \sim^0_{\mathrm{g}} \mathfrak{A}'_q, a'$ and $d(a', B') > n$ too. This is possible in sufficiently free and acyclic $\mathfrak{A}'$, since $\mathfrak{A}'_q \equiv^{\mathrm{GF}}_1 \mathfrak{A}_q$ and an n-free realisation of the appropriate type can be found according to Lemma 24.

If $d(a, B) \leqslant n$, we apply Lemma 21 to the analysis of $\hat{B} := \mathrm{cl}_n(B \cup \{a\})$. We locate the clique (guarded tuple) $\mathbf{d} \subseteq B$ in which $\hat{B} \setminus B$ is linked to B, and find its counterpart $\mathbf{d}' := \rho(\mathbf{d}) \subseteq B'$ in $\mathfrak{A}'$. As $\mathrm{cl}_n(B \cup \{a\}) = \mathrm{cl}_n(B) \cup \mathrm{cl}_n(\mathbf{d}a)$, $|\hat{B} \setminus B| \leqslant |\mathrm{cl}_n(\mathbf{d}a)| \leqslant f_n(w, w+1)$. It follows that there is a tree decomposition $\mathcal{T}$ of $\hat{B}$, in which $\mathbf{d}$ is represented at a node $v \in T$ such that the subtree $T_v \subseteq T$ that represents $\mathrm{cl}_n(\mathbf{d}a)$ has depth at most $f_n(w, w+1) \leqslant Q - q$. Now $\mathbf{d}' = \rho(\mathbf{d})$ satisfies $\varphi_{\mathcal{T},v}$ in $\mathfrak{A}'_q$, since the nesting depth of $\varphi_{\mathcal{T},v}$ is bounded by $Q - q$.

We therefore find, according to Lemma 24 and Remark 25, an extension of B' to an n-closed subset $\hat{B}'$ such that $\mathfrak{A}_q, \hat{B} \simeq \mathfrak{A}'_q, \hat{B}'$, which implies that the corresponding extension $\hat{\rho}$, as a local isomorphism over the expansions, is compatible with GF^q over $\mathfrak{A}$ and $\mathfrak{A}'$, as required. $\square$

Corollary 28. *For sufficiently large ℓ and sufficiently free and acyclic $\mathfrak{A}$ and $\mathfrak{A}'$, $\mathfrak{A} \sim^{\ell}_{\mathrm{g}} \mathfrak{A}'$ implies $\mathfrak{A} \equiv_q \mathfrak{A}'$.*

We are ready to prove expressive completeness in the sense of Theorem 26. We want to show that $\mathfrak{A} \sim^{\ell}_{\mathrm{g}} \mathfrak{A}'$ implies $\mathfrak{A} \models \varphi \;\Leftrightarrow\; \mathfrak{A}' \models \varphi$. As φ is invariant under $\sim_{\mathrm{g}}$ over finite structures, we may swap $\mathfrak{A}$ and $\mathfrak{A}'$ for structures obtained as suitable guarded bisimilar covers. This means, we can use Corollary 17 and Lemma 23 to assume w.l.o.g. that $\mathfrak{A}$ and $\mathfrak{A}'$ are conformal, N-acyclic and (n, K)-free for our preferred choice of parameters N, n, k. For some such suitable choice, Corollary 28 then implies that $\mathfrak{A} \equiv_q \mathfrak{A}'$, and $\mathfrak{A} \models \varphi \;\Leftrightarrow\; \mathfrak{A}' \models \varphi$ follows.

5. Outlook

We have introduced a new construction of finite hypergraph covers and guarded covers that seems to achieve the highest possible degree of acyclicity that can generally be guaranteed – viz., bounded acyclicity, or acyclicity in substructures of bounded size. The preliminary study of structures of bounded acyclicity has revealed some striking features, e.g., in connection with the closure operation $\mathrm{cl}_n(\,\cdot\,)$. A further model theoretic analysis of the class of structures of bounded acyclicity, including an investigation into potential algorithmic benefits of bounded acyclicity, akin in spirit maybe to that of local tree decomposability (cf. [8]), may be of interest.

The key combinatorial construction of highly acyclic Cayley groups and graphs in Section 2.1 is very uniform and the result seems natural and canonical. Its application to the cover construction is far less so, due to the local-to-global construction, in which the acyclic groups only feature as the glue. Unlike the results in the graph case [10], or the conformal covers in [9], or the new results in [3], the more highly acyclic covers obtained here are neither canonical nor compatible with automorphisms of the base structure. It remains to be seen whether this can be improved.

References

[1] N. Alon. Tools from higher algebra. In R. Graham et al., editor, *Handbook of Combinatorics*, volume II, pages 1749–1783. North-Holland, 1995.

[2] H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.

[3] V. Barany, G. Gottlob, and M. Otto. Querying the guarded fragment. This conference, 2010.

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

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

[6] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *Annals of Pure and Applied Logic*, 161:1–42, 2009. Journal version of LICS'05 paper.

[7] E. Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64:1719–1742, 1999.

[8] M. Grohe. Logic, graphs, and algorithms. In J. Flum, E. Grädel, and T. Wilke, editors, *Logic and Automata, History and Perspectives*, pages 357–422. Amsterdam University Press, 2008.

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

[10] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004. Journal version of LICS 02 paper.

[11] M. Otto. Avoiding incidental homomorphisms into guarded covers. Technical report no. 2600, TU Darmstadt, 2009.