Acyclicity in Hypergraph Covers

Martin Otto

November 16, 2009

Abstract

We construct finite groups whose Cayley graph has not only large girth, but even avoids all short cycles w.r.t. metrics induced by given subdivisions of the set of generators. Using such groups we inductively construct finite N-acyclic covers for all finite hypergraphs. As one direct application of these covers we prove that the guarded fragment has the finite model property in any class of structures defined in terms of finitely many forbidden configurations.

1 Introduction

In this technical note we deal with two main issues.

The first concerns the construction of finite Cayley groups and graphs that avoid short cycles in the strong sense that all cycles must be long even w.r.t. to a *discounted length measure*, which does not count successive steps that stay within the same subdivision of the set of generators or edge colours. This construction is entirely combinatorial; basic ideas from the well-known construction of Cayley graphs of large girth are combined with a graph-theoretic fusion of chains of Cayley graphs. This construction culminates in the construction of (rather uniform and homogeneous) Cayley graphs, for any given set of involutive generators E with a subdivision of the set E, in which every non-trivial cycle must cross boundaries between subdivisions at least N times. See Section 2 and in particular Corollary 15.

The second issue concerns the construction of finite hypergraph covers in which every configuration of up to N points is tree-decomposable (acyclic in the sense of hypergraph theory; hence, we speak of *N-acyclic covers*). This generalises corresponding constructions for graphs, where finite bisimilar covers without short cycles can be obtained as products with Cayley graphs of large girth. The generalisation to hypergraphs is non-trivial, and is here not achieved in simple products with Cayley graphs. Rather, the Cayley groups constructed in the first part are used as one stepping stone in an intricate construction of N-acyclic covers, which proceeds by induction w.r.t. the maximal cardinality of the hyperedges. The Cayley graphs are used to glue a large number of copies of an incomplete cover so as to generate a surplus of glueing sites as a supply for as yet unfulfilled requirements, and thus to complete an incomplete cover. The upshot of these considerations is the following main theorem, which will be obtained as Theorem 21 in Section 3.

Theorem. Every finite hypergraph admits, for every $N \in \mathbb{N}$, covers by finite N-acyclic and conformal hypergraphs.

Section 3.5 briefly discusses a strengthening of the finite model property for the guarded fragment, which provides one easy but strong application of the theorem, cf. Corollary 33.

2 Acyclicity in Cayley groups

2.1 Cayley groups from graphs; preliminaries

Let E be a finite set of *edge colours*, which we denote by letters e in the following. We fix some family $(E_i)_{i \in I}$ of subsets $E_i \subseteq E$ such that $E = \bigcup_i E_i$. For $\alpha \subseteq I$ we let $E_\alpha := \bigcap_{i \in \alpha} E_i$. For small subsets like $\alpha = \{i,j\}$ we also write E_{ij} for $E_{\{i,j\}} = E_i \cap E_j$; other natural relaxations with obvious meaning include abbreviations like $E_{\alpha i} = E_\alpha \cap E_i$ or $E_{\alpha\beta} = E_{\alpha \cup \beta} = E_\alpha \cap E_\beta$, etc.

We shall be dealing with E-coloured undirected graphs in which every node is incident with at most one edge of any fixed colour e, and every edge has precisely one colour (i.e., no parallel edges of distinct colours). 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 any E-coloured graph, connected components w.r.t. subsets $E' \subseteq E$ are defined as usual. Unless otherwise indicated, the E'-component is regarded as an E'-coloured graph. I.e., we pass to the reduct w.r.t. E' as well as to the restriction to the subset of nodes reachable via E'-paths from a given node. We speak of E'-components for short; in particular, an E'-component has no edges of colours $e \notin E'$. Clearly an E'-component of an E-graph is also an E'-graph.

We shall in particular look at Cayley graphs of groups generated by a set of pairwise distinct involutive generators $e \in E$. In any such group G we associate with the word $w = e_1 \cdots e_n$ over E the group element $[w]^G = e_1 \circ^G \cdots \circ^G e_n$. 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 $g \circ^G e_i$. We denote by w^{-1} the word $w^{-1} = e_n \cdots e_1$ obtained by reversing $w = e_1 \cdots e_n$; in terms of the role of these words in G, 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 an $|E|$-regular E-graph (undirected since the generators e are involutive). For subsets $E' \subseteq E$ we look at the subgroups $G \restriction E' \subseteq G$ generated by this subset of the generators of G, and at the E'-components in the Cayley graph G. Clearly, the Cayley graph of the subgroup $G \restriction E' \subseteq G$ is isomorphic to the E'-component of 1 in the Cayley graph of G.

If $v \in G$ is any node and $E' \subseteq E$, then there is a unique isomorphism of E'-coloured graphs between the E'-component of v in G and the Cayley graph of $G \restriction E' \subseteq G$ that associates v with $1 \in G \restriction E'$.

Recall the standard construction of groups with involutive generators in E from E-graphs. If H is an E-graph, we let $e \in E$ operate as an involutive permutation on the vertex set of H by letting it swap the vertices within each colour e edge of H. We denote the subgroup of the symmetric group of the vertex set of H generated by these involutive permutations as $\mathrm{Sym}(H)$.

Note that every group G with involutive generators in E is reproduced as $\mathrm{Sym}(G)$ from its own Cayley graph G. The map that sends $[w]^{\mathrm{Sym}(G)}$ to $[w]^G$ is an isomorphism, since $[w_1]^{\mathrm{Sym}(G)} \circ [w_2]^{\mathrm{Sym}(G)} = [w_1 w_2]^{\mathrm{Sym}(G)}$ just as $[w_1]^G \circ [w_2]^G = [w_1 w_2]^G$ in G.

Definition 1. Let G be a group with involutive generators E, and H any E-graph. We say that H is *compatible* with G, if for all words w over E:

$$[w]^G = 1 \quad \Rightarrow \quad [w]^{\mathrm{Sym}(H)} = 1.$$

This notion will also be applied to subgroups of the form $G \restriction E'$ for $E' \subseteq E$.

Example. Consider the group $G = \mathrm{Sym}(H)$ for some E-graph H. Then every union of connected components of H is compatible with G; similarly every union of E'-components of H is compatible with $G \restriction E'$.

If w is a word over E and $E' \subseteq E$, let $w \restriction E'$ be the word over E' obtained by deleting all letters from $E \setminus E'$ in w.

Let H be an E-graph, $H' := H \restriction E'$ the natural restriction, which only retains edges of colours from E'. Let $G = \mathrm{Sym}(H)$ and $G' = \mathrm{Sym}(H')$. Then $G' \simeq G \restriction E'$. In fact, $[w \restriction E']^G = [w]^{G'}$ because w operates like $w \restriction E'$ on H' for any w over E: all $e \in E \setminus E'$ are trivial in the absence of e-coloured edges. In particular, $[w]^G = [w]^{G'}$ for words w over E'.

Observation 2. *Suppose $G = \mathrm{Sym}(H)$, $\tilde{G} = \mathrm{Sym}(H \cup \tilde{H})$ for two disjoint E-graphs H and $\tilde{H}$. If $\tilde{H}$ is compatible with G, then $\tilde{G} \simeq G$.*

Proof. Look at the map $\pi \colon \tilde{G} \to G$ induced by the representation of group elements as sequences of generators: $\pi \colon [w]^{\tilde{G}} \mapsto [w]^G$. This is well defined, because $[w_1]^{\tilde{G}} = [w_2]^{\tilde{G}}$ implies $[w_1]^G = [w_2]^G$ because H is a connected component of $H \cup \tilde{H}$. Obviously $\pi \colon \tilde{G} \to G$ is a surjective homomorphism.

If $\tilde{H}$ is compatible with G, then this homomorphism is also injective. Suppose $[w_1]^G = [w_2]^G$; then $[w_1(w_2)^{-1}]^G = 1$ implies that $[w_1(w_2)^{-1}]^{\mathrm{Sym}(\tilde{H})} = 1$ by compatibility. Therefore w_1 and w_2 also have the same effect on any element of $\tilde{H}$. This implies that they represent the same group element also in $\tilde{G}$. □

It follows that, if $G = \mathrm{Sym}(H)$, we may assume that every E-graph that is compatible with G is represented as a disjoint connected component of H.

Definition 3. For $E' \subseteq E$, an E-graph H is called *E'-saturated* if every E'-component of H is isomorphic to some full connected component of H. I.e., for every $v \in H$, the connected component of v in the E'-reduct of H is isomorphic to a connected component of H.

Lemma 4. *Let $G = \mathrm{Sym}(H)$ for for some E-graph H, $E' \subseteq E$. Suppose that H is E'-saturated. Then for all $E'' \subseteq E$:*

$$G \restriction E' \cap G \restriction E'' = G \restriction (E' \cap E'').$$

Proof. Let $g \in G \restriction E'$ such that $g = [w]^G \in G \restriction E''$ for some word w over E''. We want to show that $g = [w \restriction E']^G$, which clearly implies membership in $G \restriction (E' \cap E'')$. It suffices to show that w and $w \restriction E'$ have the same effect on every $v \in H$. Since $g \in G \restriction E'$, the target node $v \cdot g$ lies in the E'-component of v. Considering the corresponding node v' in an isomorphic copy of this E'-component, which is a full connected component of H, we see that $v' \cdot g = v' \cdot [w \restriction E']^G$, since all $e \notin E'$ operate trivially within this component, which does not have any edges of such colours. Since the E'-components of v and v' are isomorphic, and as both g and $[w \restriction E']^G$ operate within E'-components, we find that $v \cdot g = v \cdot [w \restriction E']^G$ for all $v \in H$. Therefore $g = [w \restriction E']^G$ is an identity in G. □

Definition 5. Let $\Gamma \subseteq \mathcal{P}(E)$ be a family of subsets that is closed under intersections. We say that G *reflects intersections in* Γ if, for all $E', E'' \in \Gamma$, $G \restriction E' \cap G \restriction E'' = G \restriction (E' \cap E'')$.

Suppose that H is such that all E'-components of H are compatible with $G = \mathrm{Sym}(H)$ for all $E' \in \Gamma$. Then, by Observation 2, we may replace H by an E-graph that is E'-saturated for all $E' \in \Gamma$, without affecting G. By Lemma 4, therefore, G reflects intersections in Γ. Regarding G as $G = \mathrm{Sym}(G)$, this is in particular the case if G is compatible with (the Cayley graphs of) its restrictions (subgroups) $G \restriction E'$ for $E' \in \Gamma$.

Corollary 6. *If G is compatible with (the Cayley graphs of) its subgroups $G \restriction E'$ for all $E' \in \Gamma$, then G reflects intersections in Γ.*

2.2 Merging chains of components

Consider any two E-graphs K and K' with distinguished nodes $v \in K$ and $v' \in K'$ and a distinguished subset $\alpha \subseteq I$. Assume that the E_α-components of v and v' are isomorphic via some (in fact unique) isomorphism ρ that maps v to v'. In the situations that we shall encounter below, this will be trivially guaranteed because K and K' will be certain components of the same Cayley graph G, with E_α the intersection of the sets of edge colours in K and K'. We let

$$K \xrightarrow[\alpha]{v=v'} K'$$

be the result of glueing K and K' according to the isomorphism ρ. We say that ρ is *safe* if this merged graph is again an E-graph. This is trivially the case if E_α is the intersection of the edge colours in K and K': in this case, all edges of colours incident with both $\mathrm{dom}(\rho)$ and $\mathrm{image}(\rho)$ are covered by the isomorphism ρ.

Example 7. If K is E_i-coloured and K' is E_j-coloured then any ρ with $\alpha = \{i, j\}$ is safe. In this case, the isomorphism type of the E_i- or E_j-component of any node in the new graph is realised already in at least one of K and K'. It follows that $K \xrightarrow[\{ij\}]{v=v'} K'$ is compatible with $G \restriction E'$ for $E' = E_i, E_j$ provided K and K' were. For $\ell \neq i, j$, however, new types of E_ℓ-components may be generated in the glueing, and compatibility with $G \restriction E_\ell$ does not in general transfer.

In the following we shall build chains by merging E_α-components (isomorphic to $G_\alpha := G \restriction E_\alpha$) of the Cayley graph of a group G. In this case there always is, for any node $v \in G_\alpha$ and any α', a unique isomorphism between the $E_{\alpha\alpha'}$-components of $g \in G_\alpha$ and of $g' \in G_{\alpha'}$ (both isomorphic to $G_{\alpha\alpha'} = G \restriction E_{\alpha\alpha'}$) 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.

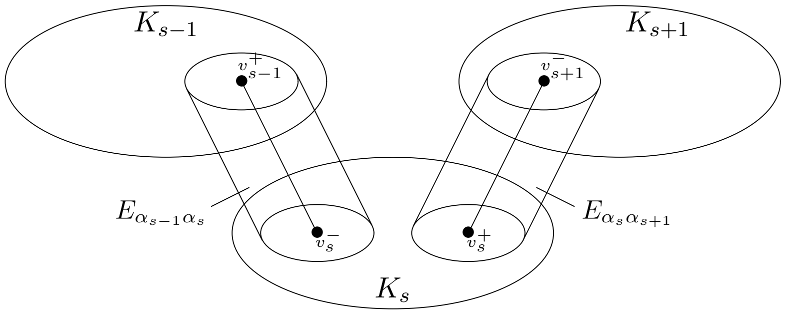

Definition 8. Consider a sequence $(K_s, v_s^-, v_s^+)_{1\leqslant s\leqslant n}$ of pairwise disjoint graphs isomorphic to E_{α_s}-components of G,

$$K_s, v_s^-, v_s^+ \simeq G_{\alpha_s}, g_s^-, g_s^+ \quad \text{for } 1 \leqslant s \leqslant n.$$[1]

This sequence is called *simple* if, for all $1 < s < n$, the connected components in K_s of v_s^- w.r.t. $E_{\alpha_{s-1}}$ and of v_s^+ w.r.t. $E_{\alpha_{s+1}}$ are disjoint.[2]

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 $E_{\alpha_{s-1}}$-component of g_s^- is disjoint from the $E_{\alpha_{s+1}}$-component of g_s^+ in G_{α_s}, or that $(g_s^-)^{-1} \circ g_s^+ \notin G_{\alpha_{s-1}\alpha_s} \circ G_{\alpha_s\alpha_{s+1}}$. It implies that the *merged chain* obtained as

$$K := \sum_{s=1}^{n} (K_s, v_s^-, v_s^+) := \quad K_1 \;\; \frac{v_1^+ = v_2^-}{\alpha_1\alpha_2} \;\; K_2 \;\; \frac{v_2^+ = v_3^-}{\alpha_2\alpha_3} \;\; \cdots \;\; \frac{v_{n-1}^+ = v_n^-}{\alpha_{n-1}\alpha_n} \;\; K_n$$

is again an E-graph.

The simplicity condition also rules out inclusion relationships between the sets of edge colours in 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 $\alpha_{s+1} \supseteq \alpha_s$ (i.e., $E_{\alpha_{s+1}} \subseteq E_{\alpha_s}$ and therefore $E_{\alpha_s\alpha_{s+1}} = E_{\alpha_{s+1}}$) rules out a continuation beyond K_{s+1}, and the merging 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 $\alpha_s = \alpha\beta_s$ for some sequence of subsets $\beta_s \subseteq I$ and a fixed subset $\alpha \subseteq I$ (α may be empty, and then we are back to the most general format).

[1] The two end points v_1^- and v_s^+ are just listed for the sake of uniformity.

[2] Weaker conditions that just require the merged sequence to be an E-graph, as in the informal notion of 'safety' above, could also be considered.

Definition 9. Let $G' \subseteq G$ be any subgroup, $\alpha \subseteq I$. We say that G' *admits chains of components* $(G_{\alpha\beta})_{\beta \subseteq I}$ *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 isomorphic to some $G_{\alpha\beta}$ for $\beta \subseteq I$.

Lemma 10. *If G_γ admits chains of length up to N of components $(G_{\alpha\beta\gamma})_{\beta \subseteq I}$ and K is a merged chain of components $(G_{\alpha\beta})_{\beta \subseteq I}$ of length up to N, then K is compatible with G_γ. In other words: if G_γ is compatible with chains of components $(G_{\alpha\beta\gamma})_{\beta \subseteq I}$ then it is also compatible with chains of components $(G_{\alpha\beta})_{\beta \subseteq I}$ up to corresponding lengths.*

2.3 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 E_i. 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 with the index set $\mathbb{Z}_n = \{0, \dots, n-1\}$ with indices modulo n.

Definition 11. Let G be a finite group with involutive generators from E, $E = \bigcup_i E_i$, with subsets $E_\alpha = \bigcap_{i \in \alpha} E_i$ and corresponding subgroups $G_\alpha = G \restriction E_\alpha$ 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 $\sigma \colon \mathbb{Z}_n \to \mathcal{P}(I) \setminus \{\emptyset\}$ such that

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

(ii) $\{g_s \colon s \in \mathbb{Z}_n\} \not\subseteq G_{\sigma(t)}$ for any single $t \in \mathbb{Z}_n$,

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

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

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.

Note that condition (iv) 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 $G_{\sigma(t)\sigma(t\pm1)} \subseteq G_{\sigma(t)}$. Intuitively, this condition says that the effect of factor g_t cannot be absorbed via variations in the immediate predecessor and successor factors.

Condition (ii) implies that σ is non-constant and thus $n \geqslant 2$. In the light of (i), condition (iv) also rules out $n = 2$ at least if G is such that $G_{\sigma(0)} \cap G_{\sigma(1)} = G_{\sigma(0)\sigma(1)}$: $g_0 \circ g_1 = 1$ implies that $g_0 = (g_1)^{-1} \in G_{\sigma(0)} \cap G_{\sigma(1)}$. The cycles we shall be interested in will always be of lengths greater than 2.

If $G_{\sigma(t)} \cap G_{\sigma(t+1)} = G_{\sigma(t)\sigma(t+1)}$, then $g_t \in G_{\sigma(t+1)}$ would already violate condition (iv): as $g_t \in G_{\sigma(t)}$ in any case, $g_t \in G_{\sigma(t+1)}$ implies that $g_t \in G_{\sigma(t)\sigma(t+1)} \subseteq G_{\sigma(t-1)\sigma(t)} \circ G_{\sigma(t)\sigma(t+1)}$. We thus have the following.

Observation 12. *If G reflects intersections in $\{E_\alpha \colon \alpha \subseteq I\}$, then condition (iv) in the definition of non-trivial coloured cycles implies condition (ii).*

2.4 Avoiding cycles of short discounted lengths

Lemma 13. *Let G be a finite group with involutive generators from $E = \bigcup_{i \in I} E_i$, $k \in \mathbb{N}$. Assume that, for all $\alpha \subseteq I$ with $|\alpha| > k$ the subgroups G_α*

(a) admit chains of components $(G_{\alpha\beta})_{\beta \subseteq I}$ 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 the same generators such that

(i) for every $\alpha \subseteq I$ with $|\alpha| > k$, $G^*_\alpha = G^* \restriction E_\alpha \simeq G \restriction E_\alpha = G_\alpha$,

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

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

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

Compare Definition 5 and Corollary 6 for the following.

Remark 14. *In the special case that $k = 0$ and for $\alpha = \emptyset$, (ii) implies in particular that G^* is compatible with its E_β-components for all $\beta \subseteq I$. Because $G^* \simeq \mathrm{Sym}(G^*)$, it follows that G^* reflects intersections in $\{E_\alpha \colon \alpha \subseteq I\}$.*

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 $|\alpha| \geqslant k$. For any such sequence, we put the corresponding merged chain

$$K = \sum_{s=1}^{n} (K_s, v_s^-, v_s^+) := \quad K_1 \; \frac{v_1^+ = v_2^-}{\alpha\beta_1\beta_2} \; K_2 \; \frac{v_2^+ = v_3^-}{\alpha\beta_2\beta_3} \; \cdots \; \frac{v_{n-1}^+ = v_n^-}{\alpha\beta_{n-1}\beta_n} \; K_n \qquad (*)$$

as a separate connected component in H^0.

By construction, $G^* = \mathrm{Sym}(G \cup H^0)$ admits chains of components $(G_{\alpha\beta})_{\beta \subseteq I}$ up to length N as required (condition (ii), first formulation). Together with (i) this implies that G^* admits chains of components $(G^*_{\alpha\beta})_{\beta \subseteq I}$ (condition (ii), second formulation) for the following reasons. If the chain in question is such that all components $G^*_{\alpha\beta}$ have $|\alpha \cup \beta| > k$, (i) tells us that $G^*_{\alpha\beta} \simeq G_{\alpha\beta}$. If on the other hand some component $G^*_{\alpha\beta}$ has $|\alpha \cup \beta| = k$, then it must be that $|\alpha| = k$ and $\beta \subseteq \alpha$ 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 $E_{\alpha'}$-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 $G^*_{\alpha'} := G^* \restriction E_{\alpha'} \simeq G_{\alpha'}$ for $|\alpha'| > k$ (cf. Observation 2).

For (iii) it remains to argue that G^*_α does not have non-trivial coloured cycles of lengths $n \leqslant N$ whenever $|\alpha| \geqslant k$. Let $|\alpha| \geqslant k$ 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 $G^*_{\sigma(t)}$ must not contain all the elements h_s according to condition (ii) of Definition 11, $G^*_{\sigma(t)} \not\supseteq G^*_\alpha$ and $\sigma(t) \not\subseteq \alpha$ for any t. It follows that $|\alpha \cup \sigma(t)| > k$. Let $h_t = [u_t]^{G^*_\alpha}$ for a word u_t over $E_{\alpha\sigma(t)}$, 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 $G_{\alpha\sigma(t)}$ will serve this purpose. We look at the sequence

$$K_s, v_s^-, v_s^+ \simeq G_{\alpha\sigma(u)}, g_s^-, g_s^+ \quad \text{with} \quad g_s^- := 1 \text{ and } g_s^+ := [u_s]^G \quad \text{for } s \in Z_n.$$

Note that, since u_s is over $E_{\alpha\sigma(s)}$, $[u_s]^G$ is the same as $[u_s]^{G_{\alpha\sigma(s)}}$, which modulo the isomorphism between $G^*_{\alpha\sigma(s)}$ and $G_{\alpha\sigma(s)}$ according to (i) is the same as $[u_s]^{G^*_{\alpha\sigma(u)}}$.

The sequence of these K_s, v_s^-, v_s^+ is simple in the sense of Definition 8, by condition (iv) in Definition 11.

It follows that $(K_s, v_s^-, v_s^-) \simeq (G_{\alpha\sigma(u_s)}, 1^G, [u_s]^G)$ is a simple sequence. 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 some value k such that conditions (a) and (b) are trivially fulfilled, we obtain the following.

Corollary 15. *Let E be a finite set decomposed into a family $E = \bigcup_{i\in I} E_i$ such that $E_\alpha := \bigcap_{i\in\alpha} E_i = \emptyset$ for all sufficiently large subsets $\alpha \subseteq I$. Then for every $N \in \mathbb{N}$ there is a finite group with E as its set of involutive generators that admits no non-trivial coloured cycles of length up to N and reflects intersections in $\{E_\alpha \colon \alpha \subseteq I\}$.*

3 Acyclic hypergraph covers

3.1 Basic definitions and preliminaries

We write $\mathfrak{A} = (A, S)$ for (finite) hypergraphs with set of hyperedges $S \subseteq \mathcal{P}(A)$. The *rank* of $\mathfrak{A}$ is the maximal cardinality of its hyperedges, denoted $\mathrm{rk}(\mathfrak{A})$.

The *Gaifman graph* associated with the hypergraph $\mathfrak{A}$ has universe A and an edge between distinct $a, a' \in A$ if $a, a' \in s$ for some hyperedge $s \in S$. Distances and *neighbourhoods* $N^L(a)$ of radius $L \in \mathbb{N}$ are defined in terms of this graph as usual.

A subset $A_0 \subseteq A$ induces the hypergraph $\mathfrak{A}\restriction A_0$ with universe A_0 and set of hyperedges $\{s \cap A_0 \colon s \in S, s \cap A_0 \neq \emptyset\}$.

Definition 16. A *cover* of a hypergraph $\mathfrak{A} = (A, S)$ is a hypergraph $\mathfrak{B} = (B, T)$ together with a surjective homomorphism $\pi \colon \mathfrak{B} \to \mathfrak{A}$ that is a hypergraph bisimulation:

(i) π maps each hyperedge of $\mathfrak{B}$ bijectively onto some hyperedge of $\mathfrak{A}$, and

(ii) satisfies the following *lifting condition*:
for all $s, s' \in S$ and every $t \in T$ with $\pi(t) = s$ there is a lift $t' \in T$ of s' at s such that $\pi(t') = s'$ and $\pi(t \cap t') = s \cap s'$.

We are interested in covers that do not have short irreducible cycles or small cliques that are not induced by single hyperedges.

Definition 17. A *cycle* in a hypergraph $\mathfrak{A} = (A, S)$ is a cycle in the graph theoretic sense in the Gaifman graph of $\mathfrak{A}$. Similarly, a *clique* in $\mathfrak{A}$ is a clique in its Gaifman graph.

A cycle of length n in $\mathfrak{A} = (A, S)$ is thus given by a cyclic word $(a_t)_{t \in \mathbb{Z}_n}$ in A such that $a_t \neq a_{t+1}$ and such that for some corresponding cyclic word $(s_t)_{t \in \mathbb{Z}_n}$ in S we have $a_t \in s_t \cap s_{t+1}$ (equivalently: $a_t, a_{t+1} \in s_{t+1}$).

A clique of size n in $\mathfrak{A} = (A, S)$ is similarly given by pairwise distinct vertices $(a_t)_{t \in \mathbb{Z}_n}$ in A such that any two of these are contained in a common hyperedge.

Definition 18. A cycle $(a_t)_{t \in \mathbb{Z}_n}$ is called *simple* if the a_t are pairwise distinct; and a simple cycle is *chordless* if no vertices $a_r, a_{r'}$ along the cycle are contained in a common hyperedge unless $r' \in \{r-1, r, r+1\}$. A cycle of length 3 is *non-trivial* if its vertices are not members of a common hyperedge. We here use the term *irreducible cycle* for chordless cycles or non-trivial 3-cycles.

Definition 19. [from classical hypergraph theory, cf. [1, 2]]

(a) A hypergraph is called *conformal* if every clique in its Gaifman graph is covered by a single hyperedge.

(b) A hypergraph is called *chordal* if every cycle of length greater than 3 in its Gaifman graph has a chord.

Finite hypergraphs cannot in general admit finite covers by chordal hypergraphs, cf. Example 31. Covers by finite conformal hypergraphs, on the other hand, are provided in [4]. The best one can hope for, in terms of chordality in finite covers, seems to be some quantitative measure of chordality that forbids at least *short* chordless cycles. We point out that the naive notion of "local chordality", viz., chordality of hypergraphs induced in Gaifman neighbourhoods of each node, cannot be achieved either, cf. Example 31.

Definition 20. A hypergraph $\mathfrak{A} = (A, S)$ is

(i) N-*chordal* if it has no irreducible cycles of length up to N.

(ii) N-*conformal* if any clique of size up to N is covered by a single hyperedge.

(iii) N-*acyclic* if the induced sub-hypergraphs on subsets $\alpha \subseteq A$ of cardinality $|\alpha| \leqslant N$ are acyclic (or tree-decomposable).

It follows from known facts in hypergraph theory [1, 2] that $\mathfrak{A}$ is N-acyclic iff it is N-conformal and N-chordal. Our goal is the following theorem.

Theorem 21. *Every finite hypergraph admits, for every $N \in \mathbb{N}$, covers by finite N-acyclic and conformal hypergraphs.*

3.2 Millefeuilles of hypergraphs

Let $\mathfrak{A} = (A, S)$ be a finite hypergraph, and let $E \subseteq S$ be a subset of the set of hyperedges, where we regard each $e \in E$ as a colour. We consider stacks of copies of the hypergraph $\mathfrak{A}$ that are joined in hyperedges $e \in E$. For $a \in A$, let $E_a := \{e \in E : a \in e\}$.

Note that the set of colours E is a set of hyperedges, while the index set I for its subdivisions is the set A of vertices of $\mathfrak{A}$.

For a group G with generator set E, we write G_a for the subgroup generated by E_a. Note that, if $a \notin \bigcup_{e \in E} e$, then $E_a = \emptyset$ and $G_a = \{1\}$.

On $A \times G$ consider the equivalence relation $\approx$ induced according to

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

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

$$[s, g] := \{[a, g] : a \in s\}.$$

For the hypergraph $\mathfrak{A} \times_E G := \hat{\mathfrak{A}}$ we put

$$\hat{\mathfrak{A}} := (\hat{A}, \hat{S}) \quad \text{where} \quad \begin{array}{rcl} \hat{A} & := & (A \times G)/{\approx}, \\ \hat{S} & := & \{[s, g] : s \in S, g \in G\}. \end{array}$$

Note that the definitions of $\approx$ and $\hat{S}$ imply

$$[a, g] \in [s, h] \quad \text{iff} \quad a \in s \text{ and } g^{-1} \circ h \in G_a \quad \text{iff} \quad a \in s \text{ and } [a, g] = [a, h].$$

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 $\mathfrak{A}$. The isomorphic copies of $\mathfrak{A}$ embedded as induced hypergraphs on these subsets $(A \times \{g\})/{\approx}$ are referred to as the *layers* of $\mathfrak{A} \times_E G$, and denoted $\mathfrak{A} \times \{g\}$. That $\mathfrak{A} \times_E G$ does not induce hyperedges on the subset $(A \times \{g\})/{\approx} \subseteq \hat{A}$ other than those of the form $[s, g]$ for $s \in S$ is clear also from the fact that the natural projection $\pi : \mathfrak{A} \times_E G \to \mathfrak{A}$ is a cover.

Recall Definitions 17 and 18 for irreducible cycles.

Proposition 22. *Let G reflect intersections in $\{E_\alpha : \alpha \subseteq A\}$, and let G have no non-trivial coloured cycles of length up to N. If $\mathfrak{A}$ does not have any irreducible cycles of length up to N, then the same is true of $\mathfrak{A} \times_E G$. In fact, any chordless cycle of length up to N in $\mathfrak{A} \times_E G$ must be contained within a single layer of $\mathfrak{A} \times_E G$, which is isomorphic to $\mathfrak{A}$.*

Proof. Let $\hat{\mathfrak{A}} = \mathfrak{A} \times_E G = (\hat{A}, \hat{S})$ and consider an irreducible cycle of length n in $\hat{\mathfrak{A}}$ given by $(\hat{a}_t)_{t \in \mathbb{Z}_n}$ in $\hat{A}$ together with $(\hat{s}_t)_{t \in \mathbb{Z}_n}$ in $\hat{S}$. This means that $\hat{s}_t = [s_t, h_t]$, $\hat{a}_t = [a_t, g_t]$ for cyclic words (s_t) over S, (a_t) over A, and (g_t) and (h_t) over G. Since $\hat{a}_t \in \hat{s}_t \cap \hat{s}_{t+1}$, it follows that $a_t \in s_t \cap s_{t+1}$ and that $(g_t)^{-1} \circ h_t \in G_{a_t}$ and $(g_t)^{-1} \circ h_{t+1} \in G_{a_t}$. In particular, $\hat{a}_t = [a_t, g_t] = [a_t, h_t]$, since $(g_t)^{-1} \circ h_t \in G_{a_t}$.

It follows that

$$u_t := (h_t)^{-1} \circ h_{t+1} = \big((g_t)^{-1} \circ h_t\big)^{-1} \circ \big((g_t)^{-1} \circ h_{t+1}\big) \in G_{a_t} \quad \text{for } t \in \mathbb{Z}_n,$$

and that $u_1 \circ \cdots \circ u_n = 1$.

Let σ be the natural colouring of $w := (u_t)_{t\in\mathbb{Z}_n}$ with $\sigma(u_t) = \{a_t\}$; in the following we often write a_t instead of $\{a_t\}$ to simplify notation. From the cyclic word $w = (u_t)_{t\in\mathbb{Z}_n}$ we want to obtain a non-trivial coloured cycle of length $\leqslant n$ in G or an irreducible cycle of length n in $\mathfrak{A}$.

The following claim helps to manipulate w and the cycle through the $(\hat{a}_t)$ so as to achieve this goal. As far as the cycle is concerned, the nodes $\hat{a}_t$ stay the same, but the hyperedges $\hat{s}_t$ joining them may get replaced. The key element in these manipulations is the replacement of representatives of the $\hat{a}_t$ so as to bring adjacent vertices and the hyperedges joining them into the same layer of $\hat{\mathfrak{A}}$. Note that for the hyperedges $\hat{s}_t = [s_t, h_t]$, $\hat{s}_t$ and $\hat{s}_{t+1}$ are in the same layer if $u_t = 1$.

Whether or not the sequence of the u_t forms a non-trivial coloured cycle, hinges on condition (iv) in Definition 11. This is because, according to Observation 12, condition (ii) in Definition 11 follows, as G reflects intersections in $\{E_\alpha \colon \alpha \subseteq A\}$.

Claim. Suppose that – contrary to condition (iv) for non-trivial coloured cycles in Definition 11 – for some factor u_t, we have $u_t \in G_{a_t a_{t-1}} \circ G_{a_t a_{t+1}}$. We want to show that either we obtain an even shorter non-trivial coloured cycle in G or the given cycle stems from an irreducible cycle in $\mathfrak{A}$.

Assume that w.l.o.g. $u_0 = k_0 \circ k_1$ for $k_0 \in G_{a_0 a_{-1}}$ and $k_1 \in G_{a_0 a_1}$.

Then $u'_{-1} := u_{-1} \circ k_0 \in G_{a_{-1}}$ and $u'_1 := k_1 \circ u_1 \in G_{a_1}$, and

$$u_{-1} \circ u_0 \circ u_1 = u'_{-1} \circ u'_1.$$

We claim that in this case, $\hat{a}_0$ is linked to $\hat{a}_{-1}$ and $\hat{a}_1$ by the hyperedges $\hat{s}'_0 := [s_0, h_0 \circ k_0]$ and $\hat{s}'_1 := [s_1, h_1 \circ (k_1)^{-1}]$, respectively. We verify that $\hat{a}_{-1}, \hat{a}_0 \in \hat{s}'_0$ and $\hat{a}_0, \hat{a}_1 \in \hat{s}'_1$:

— $\hat{a}_{-1} \in \hat{s}'_0 = [s_0, h_0 \circ k_0]$, since $(g_{-1})^{-1} \circ h_0 \circ k_0 \in G_{a_{-1}}$:
$k_0 \in G_{a_{-1}}$; and $(g_{-1})^{-1} \circ h_0 \in G_{a_{-1}}$, as $\hat{a}_{-1} \in [s_0, h_0]$.

— $\hat{a}_0 \in \hat{s}'_0 = [s_0, h_0 \circ k_0]$, since $(g_0)^{-1} \circ h_0 \circ k_0 \in G_{a_0}$:
$k_0 \in G_{a_0}$; and $(g_0)^{-1} \circ h_0 \in G_{a_0}$, as $\hat{a}_0 \in [s_0, h_0]$.

— $\hat{a}_1 \in \hat{s}'_1 = [s_1, h_1 \circ (k_1)^{-1}]$, since $(g_1)^{-1} \circ h_1 \circ (k_1)^{-1} \in G_{a_1}$:
$k_1 \in G_{a_1}$; and $(g_1)^{-1} \circ h_1 \in G_{a_1}$, as $\hat{a}_1 \in [s_1, h_1]$.

— $\hat{a}_0 \in \hat{s}'_1 = [s_1, h_1 \circ (k_1)^{-1}]$, since $(g_0)^{-1} \circ h_1 \circ (k_1)^{-1} \in G_{a_0}$:
$k_1 \in G_{a_0}$; and $(g_0)^{-1} \circ h_1 \in G_{a_0}$, as $\hat{a}_0 \in [s_1, h_1]$.

It follows that we may replace, in the original cycle, the hyperedges $\hat{s}_0$ and $\hat{s}_1$ by $\hat{s}'_0$ and $\hat{s}'_1$. We obtain a cycle linking the same nodes, with an associated coloured cycle w' in G that is shortened by one, since the factor corresponding to u_0 can be dropped. Informally, this replacement in the cyclic word $w = (u_t)$ leads to the elimination of the factor u_0 in favour of the extensions of the factors u_{-1} and u_1 which are still coloured a_{-1} and a_1 as before.

More formally, with the new hyperedges, the corresponding entry for $u_0 = (h_0)^{-1} \circ h_1$ gets replaced by

$$u'_0 = (h_0 \circ k_0)^{-1} \circ h_1 \circ (k_1)^{-1} = (k_0)^{-1} \circ (h_0)^{-1} \circ h_1 \circ (k_1)^{-1} = (k_0)^{-1} \circ k_0 \circ k_1 \circ (k_1)^{-1} = 1$$

and can be eliminated in the cyclic word w. We now assume that u_0 does not show up in w, i.e., $w = u_1 \cdots u_{n-1}$.

For the modified cycle, based on the same vertices but with the modified hyperedges as links, this means that the new hyperedges $\hat{s}'_0$ and $\hat{s}'_1$ live in the same layer of $\hat{\mathfrak{A}}$.

Since the given cycle is irreducible, $\hat{a}_{-1}$ and $\hat{a}_1$ cannot be contained in a common hyperedge. In particular, as $\hat{a}_{-1}$ and $\hat{a}_1$ live in the same layer, $a_{-1} \neq a_1$ and a_{-1} and a_1 are not members of a common hyperedge of $\mathfrak{A}$. This implies further that $E_{a_{-1}a_1} = \emptyset$ so that $G_{a_{-1}a_1} = \{1\}$.

If the next step, from $\hat{a}_1$ to $\hat{a}_2$ in the cycle, also violates condition (iv) on non-trivial coloured cycles, this now means that $u_1 \in G_{a_1a_{-1}} \circ G_{a_1a_2} = G_{a_1a_2}$, as $G_{a_1a_{-1}} = \{1\}$. If we therefore modify $\hat{s}_2 = [s_2, h_2]$ to $\hat{s}'_2 := [s_2, h_2 \circ k^{-1}]$ for $k := u_1$, we obtain instead of $u_1 = (h_1)^{-1} \circ h_2$ the new $u'_1 = (h_1)^{-1} \circ h_2 \circ k^{-1} = u_1 \circ (u_1)^{-1} = 1$. In effect, this allows us to eliminate the factor u_1 as well, and to have $\hat{s}'_2$ in the same layer with $\hat{s}'_0$ and $\hat{s}'_1$.

Either a successive elimination of factors that violate condition (iv) on non-trivial coloured cycles could eventually transform w into (a product of factors) 1: this is impossible, because that would mean that the entire cycle is represented within a single layer of the stack, and therefore would constitute an irreducible cycle in $\mathfrak{A}$; or else there remains, after a number of elimination steps, a reduced non-trivially coloured cycle in G of length up to n – which is impossible if $n \leqslant N$. ☐

The next observation will be essential towards conformality of covers.

Recall Definition 17 for cliques in a hypergraph $\mathfrak{A} = (A, S)$. Cliques in structures $\mathfrak{A} \times_E G$ as constructed above, project injectively onto cliques in $\mathfrak{A}$, since distinct nodes $\hat{a} = [a, g] \neq [a, g'] = \hat{a}'$ above the same node $a \in A$ cannot be linked by any hyperedge. It follows that for conformal $\mathfrak{A}$, $\mathfrak{A} \times_E G$ cannot have any cliques of size greater than $\mathrm{rk}(\mathfrak{A})$. Note that cliques of size 3 are the same as cycles of length 3.

Proposition 23. *Let G reflect intersections in $\{E_\alpha \colon \alpha \subseteq A\}$, and let G have no non-trivial coloured cycles of length up to N. Then any clique of size up to N in $\mathfrak{A} \times_E G$ is fully contained in some single layer isomorphic to $\mathfrak{A}$.*

Proof. Let n be minimal such that $\hat{\mathfrak{A}} := \mathfrak{A} \times_E G$ has a clique of size n *not* contained in a single layer. Let $(\hat{a}_t)_{t \in \mathbb{Z}_n}$ be such a clique, where $\hat{a}_t = [a_t, g_t]$ for some clique $(a_t)_{t \in \mathbb{Z}_n}$ and suitable $g_t \in G$. Let $\alpha = \{a_t \colon t \in \mathbb{Z}_n\}$. By minimality of n we have that every subset of up to $n - 1$ elements among the $\hat{a}_t$ is represented within a single layer of the stack. In particular, for every $t \in \mathbb{Z}_n$, there is some $h_t \in G$ such that $[a_r, g_r] = [a_r, h_t]$ for all $r \neq t$. Consider then the group elements

$$u_t := (h_t)^{-1} \circ h_{t+1} \quad \text{for } t \in \mathbb{Z}_n.$$

Clearly $u_1 \circ \cdots \circ u_n = 1$.

By our assumptions, $(g_r)^{-1} \circ h_t \in G_{a_r}$ for all $r \neq t$, and therefore $(h_t)^{-1} \circ h_{t+1} = ((g_r)^{-1} \circ h_t)^{-1} \circ ((g_r)^{-1} \circ h_{t+1}) \in G_{a_r}$ for all $r \neq t, t+1$.

It follows that $u_t \in G_{\alpha_t}$ for $\alpha_t := \alpha \setminus \{a_t, a_{t+1}\}$ and that $\sigma \colon \mathbb{Z}_n \to \mathcal{P}(\alpha) \setminus \{\emptyset\}$, $\sigma(u_t) := \alpha_t$ is a colouring.

We claim that, since $(\hat{a}_t)_{t \in \mathbb{Z}_n}$ is not contained in any single layer, $(u_t)_{t \in \mathbb{Z}_n}$ with the colouring σ yields a non-trivial coloured cycle in G, whence $n \geqslant N$ follows.

We need to verify condition (iv) of Definition 11, i.e., that $u_t \notin G_{\alpha_{t-1}\alpha_t} \circ G_{\alpha_t\alpha_{t+1}}$. Note that $\alpha_{t-1} \cup \alpha_t = \alpha \setminus \{a_t\}$ and $\alpha_t \cup \alpha_{t+1} = \alpha \setminus \{a_{t+1}\}$.

Reasoning towards a contradiction, suppose that $u_t = k_0 \circ k_1$ with $k_0 \in G_{\alpha_{t-1}\alpha_t}$ and $k_1 \in G_{\alpha_t\alpha_{t+1}}$. Then

$$\begin{aligned} \hat{a}_r &= [a_r, h_t] = [a_r, h_t \circ k_0] && \text{for } r \neq t, \text{ and} \\ \hat{a}_r &= [a_r, h_{t+1}] = [a_r, h_{t+1} \circ (k_1)^{-1}] && \text{for } r \neq t+1. \end{aligned}$$

For instance, for the first of these assertion note that $k_0 = (h_t)^{-1} \circ h_t \circ k_0 \in G_{\alpha_{t-1}\alpha_t} \subseteq G_{a_r}$ for all $r \neq t$. But

$$h_{t+1} \circ (k_1)^{-1} = h_t \circ (h_t)^{-1} \circ h_{t+1} \circ (k_1)^{-1} = h_t \circ u_t \circ (k_1)^{-1} = h_t \circ k_0.$$

It follows that in this case $\hat{a}_r = [a_r, h_t \circ k_0]$ for $r \neq t$ and $\hat{a}_t = [a_t, h_{t+1}] = [a_t, h_{t+1} \circ (k_1)^{-1}] = [a_t, h_t \circ k_0]$ are all represented in the same layer. This contradicts the asumptions about $(\hat{a}_t)_{t \in \mathbb{Z}_n}$.

$\square$

Corollary 24. *Let G reflect intersections in $\{E_\alpha \colon \alpha \subseteq A\}$, and let G have no non-trivial coloured cycles of length up to N. Then any irreducible cycle of length up to N as well as any clique of size up to N in $\mathfrak{A} \times_E G$ is fully contained in a single layer of $\mathfrak{A} \times_E G$ isomorphic to $\mathfrak{A}$.*

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

If $\mathfrak{A}$ is conformal and $N \geqslant \mathrm{rk}(\mathfrak{A})$, then $\mathfrak{A} \times_E G$ is also conformal.

Suppose $\mathfrak{A} = (A, S)$ and $s_0 \in S$ and $E \subseteq S$ are such that the (Gaifman) distance from s_0 to any $s \in E$ is at least N.

Think of s_0 as in the centre of $\mathfrak{A}$ while $E \subseteq S$ consists of hyperedges in the periphery. We shall later use this in the situation where an interior part of $\mathfrak{A}$ already behaves like an acyclic cover but is incomplete in this respect in its boundary region; missing hyperedge neighbours of peripheral hyperedges in this boundary region will be supplied through glueing with central hyperedges in new copies of $\mathfrak{A}$. For this we need a surplus of central hyperedges in $\mathfrak{A}$ compared to the demands created by its peripheral hyperedges, and it is to this end that stacking is used: to create many layers of copies of interior hyperedges without unduly increasing the number of peripheral ones, as follows.

In the given situation, the glueing of isomorphic copies of $\mathfrak{A}$ produces a hypergraph $\mathfrak{A}' = \mathfrak{A} \times_E G$. This new hypergraph $\mathfrak{A}' = (A', S')$ is a cover of $\mathfrak{A}$ w.r.t. the natural projection $\pi \colon \mathfrak{A}' \to \mathfrak{A}$, such that any one of the copies of s_0 in the different layers of $\mathfrak{A}'$ are far from each other and far from the copies of elements $s \in E$. Moreover, the multiplicity ratio between center and boundary is improved by at least a factor of two: $\mathfrak{A}'$ has $|G|$ many disjoint isomorphic copies of s_0 above s_0, and at most $|G|/2$ many distinct copies of any $s \in E$ above s:

for copies of s_0 this is because the equivalence classes of $[a, g]$ for $a \in s_0 \subseteq A \setminus \bigcup_{s \in E} s$ are singletons;

for copies of $s \in E$, we have a factor of at least two, because $[s, g] = [s, g \circ e]$ for the generator $e = s$ of G, and $g \neq g \circ e$.

If we now let $E' \subseteq S'$ consist of all the hyperedges above the $s \in E$, then the situation is as before and we may repeat the process of stacking with a group G' that has generators corresponding to the hyperedges in E'.

Let, for a cover $\pi\colon \mathfrak{A}' \to \mathfrak{A}$ and any $s \in S$, the *multiplicity of π over s* be defined as the cardinality of the fiber above s:

$$\mu(\pi, s) := |\{s' \in S' \colon \pi(s') = s\}|.$$

Repeated application of stacking with glueing in copies above the original hyperedges in E leads to a series of hypergraph covers $\pi_n\colon \mathfrak{A}_n \to \mathfrak{A}$ where $\mathfrak{A}_n = (A_n, S_n)$, starting with $\mathfrak{A}_0 = \mathfrak{A} = (A, S)$, and such that

$$\mathfrak{A}_{n+1} = \mathfrak{A}_n \times_{E_n} G_n$$

where

(i) $E_n \subseteq S_n$ consists of the set of hyperedges that are mapped onto those in $E \subseteq S$ by π_n, and

(ii) G_n is a group with generators corresponding to the hyperedges in E_n.

Then, in this sequence, the ratio

$$\frac{\mu(\pi_n, s_0)}{\mu(\pi_n, s)} \geqslant 2^n \qquad \text{for any } s \in E$$

becomes arbitrarily large. If $\mathfrak{A}$ is N-acyclic and the G_n are chosen without short non-trivial coloured cycles and reflecting intersections, then each cover $\pi_n\colon \mathfrak{A}_n \to \mathfrak{A}$ remains N-acyclic (and conformal if $\mathfrak{A}$ is conformal and $N \geqslant \mathrm{rk}(\mathfrak{A})$).

The same applies to any set of central hyperedges E_0 that are far from the peripheral hyperedges in E and we obtain the following corollary.

Corollary 25. *Every finite N-acyclic hypergraph $\mathfrak{A} = (A, S)$ with subsets $E_0, E \subseteq S$ at distance greater than N (i.e., with $d(a, a') > N$ for every $a \in s \in E_0$ and $a' \in s' \in E$) admits, for every $m \in \mathbb{N}$ a cover*

$$\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$$

by a finite N-acyclic hypergraph $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ such that

(i) $\hat{E}_0 := \{\hat{s} \in \hat{S} \colon \pi(\hat{s}) \in E_0\}$ has distance greater N from $\hat{E} := \{\hat{s} \in \hat{S} \colon \pi(\hat{s}) \in E\}$;

(ii) any two distinct $\hat{s}, \hat{s}' \in \hat{S}$ above the same $s = \pi(\hat{s}) = \pi(\hat{s}') \in E_0$ have distance greater than N (in fact even greater than $2N$);

(iii) the multiplicity ratio $\mu(\pi, s_0)/\mu(\pi, s)$ between hyperedges above $s_0 \in E_0$ and hyperedges above $s \in E$ is at least m.

If $\mathfrak{A}$ is conformal and $N \geqslant \mathrm{rk}(\mathfrak{A})$, then $\hat{\mathfrak{A}}$ is conformal, too.

3.3 Local covers

In a first step we want to obtain L-local N-acyclic covers at every a. The construction of these will rely on the availability of (full rather than local) N-acyclic and conformal covers of hypergraphs whose hyperedges are smaller than those of $\mathfrak{A}$. Recall that the *rank* of a hypergraph $\mathfrak{A} = (A, S)$, $\mathrm{rk}(\mathfrak{A})$, is the maximal cardinality of its hyperedges. The construction of local covers for $\mathfrak{A}$ uses full covers for certain derived hypergraphs of rank $\mathrm{rk}(\mathfrak{A}) - 1$. The basic step in the construction is reflected in the following simple observations.

For technical reasons we may assume that the set of hyperedges is closed under subsets, i.e., $s' \subseteq s \in S$ implies $s' \in S$.

Consider a node a in a hypergraph $\mathfrak{A} = (A, S)$. The *localisation* of $\mathfrak{A}$ at a is the hypergraph $\mathfrak{A} \restriction N^1_*(a)$ 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 $N^1_*(a)$.

Observation 26. *If $\mathfrak{A}$ is conformal, then the rank of any localisation of $\mathfrak{A}$ is strictly less than the rank of $\mathfrak{A}$.*

Observation 27. *Let $a \in \mathfrak{A} = (A, S)$ and $\pi\colon \mathfrak{B}_0 \to \mathfrak{A} \restriction N^1_*(a)$ with $\mathfrak{B}_0 = (B_0, T_0)$ a cover. Then, for a new element $b \notin B_0$, the hypergraph $\mathfrak{B} := (B, T)$ with $B = B_0 \cup \{b\}$ and*

$$T = \{t \in T_0 \colon \pi(t) \in S \restriction N^1(a)\} \cup \{t \cup \{b\} \colon t \in T_0, \pi(t) \cup \{a\} \in S \restriction N^1(a)\}$$

with the natural extension of π, which maps b to $\pi(b) := a$, provides a cover of $\mathfrak{A} \restriction N^1(a)$ at a. Moreover,

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

(ii) if $\mathfrak{B}_0$ is N-chordal and $\mathfrak{A}$ is (at least 3-)conformal, then so is $\mathfrak{B}$.

Proof. For (i) consider cliques in $\mathfrak{B}$. If the clique is contained in B_0, (N-)conformality of $\mathfrak{B}_0$ settles this. A clique including $b \in B$ must be of the form $t \cup \{b\}$ for a clique $t \subseteq B_0$ which therefore is a hyperedge $t \in T_0$; but then $\pi(t) \cup \{a\}$ is a clique in $\mathfrak{A}$ and thus in $S \restriction N^1(a)$, and hence $t \cup \{b\}$ was turned into a hyperedge of $\mathfrak{B}$.

For (ii), similarly, the case of cycles with nodes just from B_0 is settled in $\mathfrak{B}_0$; and any cycle involving $b \in B$ of length greater than 3 is chordal as any node of $\mathfrak{B}$ is linked to b by a hyperedge (for cycles of length 3 involving b, we argue as for 3-cliques and use 3-conformality of $\mathfrak{A}$). $\square$

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

Lemma 28. *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 $\mathfrak{B}$.

Proof. $\mathfrak{B}$ is obtained by glueing one new disjoint isomorphic copy of $\mathfrak{C}$ onto each individual hyperedge of $\mathfrak{B}_0$.

Consider a hyperedge t of $\mathfrak{B}_0$ with image $s = \pi_0(t)$ in $\mathfrak{A}$. Let $\rho^{(t)}\colon \mathfrak{C}^{(t)} \to \mathfrak{A}$ be a fresh isomorphic copy of the cover $\rho\colon \mathfrak{C} \to \mathfrak{A}$. In $\mathfrak{C}^{(t)}$, choose a hyperedge $t' \subseteq C^{(t)}$ above $\pi_0(t) = s$. Let $f^{(t)}\colon t' \to t$ be the bijection between $t' \subseteq C^{(t)}$ and $t \subseteq B_0$, which is induced by $\rho^{(t)}$ and π_0, i.e., such that $\pi_0 \circ f^{(t)} = \rho^{(t)} \restriction t'$.

We let $\mathfrak{B}$ be the hypergraph obtained by glueing $\mathfrak{B}_0$ and all the disjoint $\mathfrak{C}^{(t)}$, where each $\mathfrak{C}^{(t)}$ is glued via the corresponding $f^{(t)}$, so as to identify just the chosen $t' \subseteq C^{(t)}$ and $t \subseteq B_0$.

It is clear that $\mathfrak{B}_0 \subseteq \mathfrak{B}$ and that $\pi \colon \mathfrak{B} \to \mathfrak{A}$ is a cover of the required kind. Moreover, (N-)conformality and N-chordality are preserved in this glueing:

(i) every clique in $\mathfrak{B}$ is fully contained in B_0 or in one of the $\mathfrak{C}^{(t)}$.

(ii) every irreducible cycle is fully contained in B_0 or in one of the $\mathfrak{C}^{(t)}$.

For the second claim, consider a cycle linking nodes in $B_0 \setminus t$ to nodes in $C^{(t)} \setminus t'$; as the identification of t with t' is the only bridge between these two parts, the cycle would have to pass through this common patch at least twice; as this common part is a hyperedge of $\mathfrak{B}$, this induces a chord and the cycle cannot be irreducible. □

Definition 29. Let $\mathfrak{A} = (A, S)$ and $\mathfrak{B} = (B, T)$ be hypergraphs with distinguished nodes $a \in A$ and $b \in B$, $L \in \mathbb{N}$. Then $\pi \colon \mathfrak{B}, b \to \mathfrak{A}, a$ is called an *L-local cover* if

(i) π maps every hyperedge of $\mathfrak{B}$ bijectively onto a hyperedge of $\mathfrak{A}$, and $\pi(b) = a$.

(ii) the lifting condition of Definition 16 is satisfied for all $s, s' \in S$ with $s \cap s' \neq \emptyset$ and every $t \in T$ with $\pi(t) = s$ for which $t \subseteq N^{L-1}(b) \subseteq B$.[3]

Lemma 30. *Suppose that N-acyclic, conformal covers are available for all rank k hypergraphs. Then there is, for every hypergraph $\mathfrak{A}$ of rank $k+1$, every element $a \in A$ and every $L \in \mathbb{N}$, an L-local cover $\pi \colon \mathfrak{B}, b \to \mathfrak{A}, a$ by an N-acyclic and conformal hypergraph $\mathfrak{B}$.*

Proof. The construction of $\pi \colon \mathfrak{B}, b \to \mathfrak{A}, a$ is by induction on the radius L, starting from a cover of the localisation $\mathfrak{A} \restriction N^1_*(a)$ and of $\mathfrak{A} \restriction N^1(a)$ (as in Observation 27). We successively extend incomplete 1-neighbourhoods of points $b' \in N^{L-1}(b)$ to 1-neighbourhoods that provide covers for $\mathfrak{A} \restriction N^1(\pi(b'))$. Let $\mathfrak{B}_0$ be the current, incomplete N-acyclic cover, b' on the boundary in the sense that $\mathfrak{B}_0 \restriction N^1(b')$ is not yet a full cover of $\mathfrak{A} \restriction N^1(\pi(b'))$.

The extension step is performed at the level of rank k hypergraphs:

– we extend the partial cover of $\mathfrak{A} \restriction N^1_*(\pi(b'))$ provided by $\mathfrak{B}_0 \restriction N^1_*(b')$ to a full cover of $\mathfrak{A} \restriction N^1_*(\pi(b'))$ according to Lemma 28,

– we fill in b' according to the trick in Observation 27, to obtain a full cover $\mathfrak{B}^1, b'$ of $\mathfrak{A} \restriction N^1(\pi(b'))$ that has $\mathfrak{B}_0 \restriction N^1(b')$ as a substructure, and

– we glue this cover $\mathfrak{B}^1, b'$ to $\mathfrak{B}_0$ in $\mathfrak{B}_0 \restriction N^1(b')$ (this part is common to both hypergraphs, which are taken to be otherwise disjoint).

(N-)conformality and N-chordality are preserved in this glueing as well.

This is clear for conformality: any clique in the resulting structure must be contained in either of the two parts, as no new (hyper-)edges are introduced.

For N-chordality consider a chordless cycle in the resulting structure that is not fully contained in either of the two parts, $\mathfrak{B}_0$ or $\mathfrak{B}^1, b'$. Since the cycle is not contained in $\mathfrak{B}^1, b'$, it must have at least two nodes at distance greater than 1 in $\mathfrak{B}_0 \restriction N^1_*(b')$, that are linked by a segment of the cycle that is fully within $\mathfrak{B}_0$. If $b_1, b_2 \in \mathfrak{B}_0 \restriction N^1_*(b')$ are such, then we may close this segment to form a new cycle by filling in b' between b_1 and b_2. This cycle would be chordless in $\mathfrak{B}_0$, and can only be shorter than the given one; hence the given one had length greater than N. □

[3]This means that the corresponding lift t' of s' at t will still be contained in $N^L(b)$.

Availability of N-acyclic covers for rank 2 hypergraphs follows from [5]. Note that rank 2 hypergraphs are graphs $\mathfrak{A} = (A, E)$, and the basic construction of Cayley groups from graphs (in this case, 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 rank.

3.4 From local to global covers through stacking and glueing

Let $\mathfrak{A} = (A, S)$ be any finite hypergraph. Using the construction from [4], $\mathfrak{A}$ has a finite conformal cover. Let us therefore assume w.l.o.g. that $\mathfrak{A}$ itself is conformal. We assume inductively that all hypergraphs of rank less than $\mathrm{rk}(\mathfrak{A})$ admit N-acyclic covers.

We also assume w.l.o.g. that the Gaifman graph of $\mathfrak{A}$ is connected and let $a \in A$ and $r \in \mathbb{N}$ be such that $A \subseteq N^r(a)$.

With Lemma 30 we obtain an $(N + r + 2)$-local N-acyclic cover $\pi_0\colon \mathfrak{B}, b \to \mathfrak{A}, a$.

If $\mathfrak{B} = (B, T)$, it follows that for every $s \in S$ there is some $t \in T$ within $N^r(b)$ above s, i.e., such that $\pi_0(t) = s$. Let $E_0 \subseteq T$ the set of hyperedges within $N^r(b)$, and $E \subseteq T$ the set of hyperedges of $\mathfrak{B}$ that are *not* contained within $N^{N+r+1}(b)$ (it is for these that $\mathfrak{B}$ may not have hyperedge neighbours to lift hyperedge neighbours from $\mathfrak{A}$).

It follows from these choices that the distance between E_0 and E is greater than N.

So Corollary 25 applies to yield N-acyclic covers for $\mathfrak{B}$ such that hyperedges above hyperedges from E_0 are at distance greater than N from those above E and from each other, and such that the multiplicity ratio between hyperedges above E_0 and those above E is as large as desired. Let $m := |E|$ and choose $\pi\colon \hat{\mathfrak{B}} \to \mathfrak{B}$ such that $\mu(\pi, t_0)/\mu(\pi, t) \geqslant m$ for every $t_0 \in E_0$ and $t \in E$.

Consider the set of critical hyperedges $\hat{E} := \{\hat{t} \in \hat{T}\colon \pi(\hat{t}) \in E\} \subseteq \hat{T}$ of $\hat{\mathfrak{B}} = (\hat{B}, \hat{T})$, which may be short of hyperedge neighbours. For its cardinality we have

$$|\hat{E}| \leqslant \mu(\pi, t)|E| \leqslant \mu(\pi, t_0) =: \mu$$

for any $t \in E$ and $t_0 \in E_0$, where μ is the interior multiplicity. This μ is the number of layers whose disjoint interiors are all isomorphic to $\mathfrak{B} \restriction N^r(b)$ in $\hat{\mathfrak{B}}$. We also write $\hat{E}_0 := \{\hat{t_0} \in \hat{T}\colon \pi(\hat{t_0}) \in E_0\} \subseteq \hat{T}$ for the hyperedges in these μ many disjoint interior parts of layers of $\hat{B}$ above E_0.

We may now satisfy all requirements for hyperedge neighbours at all critical hyperedges $\hat{t} \in \hat{E}$ by identifying (glueing) $\hat{t}$ with suitable partners $f(\hat{t}) \in \hat{E}_0$ in pairwise distinct interior layers. For this, choose f_0 as an injection of $\hat{E}$ into $\mu = \{0, \ldots, \mu - 1\}$ and pick, for each $\hat{t} \in \hat{E}$ some partner hyperedge $f(\hat{t})$, in the layer given by $f_0(\hat{t})$, for which $\pi_0(\pi(\hat{t})) = \pi_0(\pi(f(\hat{t})))$ (equality in the set S of hyperedges of $\mathfrak{A}$).

Let $\hat{\mathfrak{A}} = (\hat{A}, \hat{S})$ be the resulting hypergraph, with node set $\hat{A}$ obtained from $\hat{B}$ by identification of nodes in $\hat{t}$ with those in $f(\hat{t})$ for $\hat{t} \in \hat{E}$, and hyperedge set $\hat{S}$ induced by $\hat{T}$. Let $\hat{\pi}\colon \hat{A} \to A$ be the projection induced on this quotient by $\pi_0 \circ \pi$ (the choice of f is such that the quotient is compatible with $\pi_0 \circ \pi$).

We claim that $\pi\colon \hat{\mathfrak{A}} \to \mathfrak{A}$ is a cover and that $\hat{\mathfrak{A}}$ is N-acyclic and conformal. Both claims are in fact obvious from the construction.

For the cover property: lifts at hyperedges $\hat{t} \notin \hat{E}$ were even guaranteed in $\hat{\mathfrak{B}}$; for lifts at hyperedges $\hat{t} \in \hat{E}$, the identification of these hyperedges with interior hyperedges means that they inherit all required neighbours from those.

For N-acyclicity: $\hat{\mathfrak{B}}$ is N-acyclic and we just need to convince ourselves that the glueing via f does not disturb N-chordality or N-conformality. Because the identifications according to f are with pairwise far apart interior hyperedges, it is clear that a connected configuration of up to N many points cannot connect nodes in the interior of a layer to the exterior of another layer other than through f-identifications. It follows that any such configuration is isomorphic to one realised in some hypergraph obtained by glueing a finite number of disjoint copies of interior patches of individual layers of $\hat{\mathfrak{B}}$ onto exterior hyperedges in $\hat{\mathfrak{B}}$, as with f. In this process of glueing N-acyclic hypergraphs in single hyperedges, N-acyclicity is preserved, and so is conformality (compare the argument in connection with Lemma 28).

We have thus proved Theorem 21, i.e., that every finite hypergraph admits, for every $N \in \mathbb{N}$, covers by finite N-acyclic and conformal hypergraphs. While we thus achieve unqualified conformality in our finite covers, just as in [4], it is clear that some qualification is necessary w.r.t. chordality.

Example 31. Consider the cycle graph $\mathfrak{C}_n$ with vertex set $C_n := \mathbb{Z}_n$ and edge set $R = \{\{t, t+1\}\colon t \in \mathbb{Z}_n\}$. Its rank 3 hypergraph companion is obtained by the addition of a new central node a and extension of edges to hyperedges through joining with a:

$$\mathfrak{A} = (A, S) \quad \text{with } A = C_n \,\dot{\cup}\, \{a\} \text{ and } S = \{\{t, t+1, a\}\colon t \in \mathbb{Z}_n\}.$$

Note that $\mathfrak{C}_n$ is the localisation of $\mathfrak{A}$ at a, i.e. its restriction to $N^1_*(a)$ (cf. the construction and discussion in connection with Observation 27).

If $\pi\colon \mathfrak{B} \to \mathfrak{A}$ is any cover, then the restriction of π to the localisation of $\mathfrak{B}$ at any $b \in \pi^{-1}(a)$, i.e., its restriction to $N^1_*(b)$, will be a cover of $\mathfrak{C}_n$.

The graph $\mathfrak{C}_n$ clearly does not have finite acyclic covers. Any cycle of the cover of $\mathfrak{C}_n$ induced by a restriction of π to some $N^1_*(b)$ comes from a chordless cycle in $\mathfrak{B}$ and in $\mathfrak{B} \restriction N^1(b)$. It follows that $\mathfrak{A}$ cannot have any finite chordal covers; in fact not even finite covers in which the 1-neighbourhoods of nodes are chordal.

3.5 Application to a finite model property of GF

One immediate corollary concerns a substantial strengthening of the finite model property for the guarded fragment GF (cf. [3, 4]) to a finite model property over $\mathcal{C}$ for any class $\mathcal{C}$ defined in terms of forbidden cyclic configurations.

Fix a finite relational vocabulary τ.

Definition 32. A class $\mathcal{C}$ of (finite and infinite) τ-structures *excludes the configuration* $\mathfrak{A}_0$, if no substructure of any $\mathfrak{A} \in \mathcal{C}$ is isomorphic to $\mathfrak{A}_0$.

A class $\mathcal{C}$ of (finite and infinite) τ-structures is *defined by finitely many excluded configurations* if, for some finite collection of finite τ-structures $(\mathfrak{A}_i)$, $\mathcal{C}$ consists of those τ-structures that do not have any structure isomorphic to one of the $\mathfrak{A}_i$ as a finite substructure.

Note that it does not matter whether we speak of finitely many forbidden configurations or of configurations of some bounded size that are ruled out. One could therefore also define such a class in terms of a positive list of allowed configurations up to some threshold size that may occur as substructures – i.e. in terms of a finite initial segment of its *age*.

A logic has the finite model property over $\mathcal{C}$ if every formula that is satisfiable in $\mathcal{C}$ has a finite model in $\mathcal{C}$. A combination of the generalised tree model property of GF (see [3], but here we really appeal to the guarded-tree-decomposable model property) and our finite covers (which directly also produce N-acyclic guarded bisimilar covers of finite relational structures), then yields the following.

Corollary 33. *Let $\mathcal{C}$ be a class of τ-structures defined in terms of some finite collection $(\mathfrak{A}_i)_{i \leqslant N}$ of finite excluded configurations where each $\mathfrak{A}_i$ is* cyclic *in the sense of not being guarded tree-decomposable. Then* GF *has the finite model property over $\mathcal{C}$.*

Note the important constraint here, that the forbidden substructures $\mathfrak{A}_i$ are cyclic. It is not possible (not with our techniques, but also provably impossible) to exclude certain acyclic configurations. For instance, if it were possible to exclude *all* expansions of the three-node directed E-tree with a single root and two sibling nodes attached by E-edges, one would obtain a finite model property for GF with number constraints of the form $\exists^{\leqslant 1} yExy$. But this extension of GF is known to be an undecidable fragment of FO, [3], hence cannot have a finite model property. Similarly, at least our techniques cannot exclude guarded tree-decomposable embeddings of short cycles. For instance, consider a partial duplication (in a cover) of a ternary hyperedge $\{a, b, c\}$ to two overlapping hyperedges $\{a, b, c\}$ and $\{a, b, c'\}$, which yields the 4-cycle a, c, b, c', albeit in a guarded tree-decomposable embedding.

Thanks to Achim Blumensath for a careful reading of an earlier draft of this report.

References

[1] Catriel Beeri, Ronald Fagin, David Maier, and Mihalis Yannakakis, *On the desirability of acyclic database schemes*, Journal of the ACM 30, pp. 479-513, 1983.

[2] Claude Berge, Graphs and Hypergraphs, North-Holland, 1973.

[3] Erich Grädel, *The restraining power of guards*, Journal of Symbolic Logic 64, pp. 1719–1742, 1999.

[4] Ian Hodkinson and Martin Otto, *Finite conformal hypergraph covers and Gaifman cliques in finite structures*, The Bulletin of Symbolic Logic 9, 2003, pp. 387–405.

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