THE BULLETIN OF SYMBOLIC LOGIC
Volume 9, Number 3, Sept. 2003

FINITE CONFORMAL HYPERGRAPH COVERS AND GAIFMAN CLIQUES IN FINITE STRUCTURES

IAN HODKINSON AND MARTIN OTTO

Abstract. We provide a canonical construction of conformal covers for finite hypergraphs and present two immediate applications to the finite model theory of relational structures. In the setting of relational structures, conformal covers serve to construct guarded bisimilar companion structures that avoid all incidental Gaifman cliques—thus serving as a partial analogue in finite model theory for the usually infinite guarded unravellings. In hypergraph theoretic terms, we show that every finite hypergraph admits a bisimilar cover by a finite conformal hypergraph. In terms of relational structures, we show that every finite relational structure admits a guarded bisimilar cover by a finite structure whose Gaifman cliques are guarded. One of our applications answers an open question about a clique constrained strengthening of the extension property for partial automorphisms (EPPA) of Hrushovski, Herwig and Lascar. A second application provides an alternative proof of the finite model property (FMP) for the clique guarded fragment of first-order logic CGF, by reducing (finite) satisfiability in CGF to (finite) satisfiability in the guarded fragment, GF.

§0. Introduction. The main construction in this paper is presented in terms of hypergraphs, i.e., structures consisting of just a universe together with a collection of subsets of the universe. The main motivation behind the construction, however, arises in the context of ordinary relational structures; and here primarily their model theory with respect to guarded logics and extension properties for partial automorphisms.

Guarded logics, as introduced by Andréka, van Benthem and Németi in [1], play an important role in transferring some of the good algorithmic properties of modal logics to the setting of general relational structures. In particular, they provide natural decidable fragments of first-order logic at an interesting level of expressiveness. The guarded fragment GF of first-order

Received August 27, 2002; revised February 10, 2003.

2000 *Mathematics Subject Classification.* primary 03C13, secondary 03B45, 03B70, 05C65, 05C69.

Key words and phrases. finite model theory, extension property for partial isomorphisms, guarded logics, finite model property.

Ian Hodkinson's research partially supported by EPSRC grant GR/R45369/01.

Martin Otto's research partially supported by EPSRC grant GR/R11896/01.

The authors thank the referee for helpful suggestions.

1079-8986/03/0903-0005/$2.90

logic, and its further generalisation to the clique guarded fragment CGF, capture relativised modes of quantification that restrict the access to certain configurations of elements in a structure. For GF, only guarded configurations, consisting of elements that coexist within some relational ground atom (their guard), are simultaneously accessible. For CGF, more general configurations consisting of elements that are pairwise guarded (cliques in the Gaifman graph, see below) are accessible.

Just as bisimulation equivalence is the fundamental invariance for modal logic, there is a corresponding notion of *guarded bisimulation equivalence* which provides the crucial invariance for GF. The study of guarded bisimulation invariance gives rise to tree-like guarded bisimilar companion structures—indistinguishable from the given structure in GF but coming with a natural tree decomposition. They are obtained through a process of guarded unravelling, analogous to bisimilar unravellings of graphs or transition systems into actual trees. These tree-like companions feature prominently in many model theoretic arguments for GF; compare, for instance, [9, 7, 10]. But guarded unravellings, even of finite structures, are infinite in general. They therefore do not usually lend themselves to arguments in the finite model theory of guarded logics.

Recall that the *Gaifman graph* $G(\mathfrak{A})$ of a relational structure $\mathfrak{A} = (A, \dots)$ has for its vertex set the universe A of $\mathfrak{A}$, and edges between any two distinct elements that occur together in a common relational ground atom of $\mathfrak{A}$. We are here interested in cliques in the Gaifman graph, or Gaifman clique for short. A *Gaifman clique* in $\mathfrak{A}$ is formed by any subset of A such that any two distinct elements inside this subset are linked by an edge of $G(\mathfrak{A})$, i.e., occur together in a ground atom of $\mathfrak{A}$. There are two motivations for looking at Gaifman cliques. From the point of view of guarded logics, Gaifman cliques are precisely the configurations accessible in the clique guarded fragment CGF. More importantly, Gaifman cliques arise as natural obstacles for tree decompositions of relational structures—in fact certain kinds of Gaifman clique, besides certain kinds of cycle, are the only types of obstacle (see section 1.2 below).

There are two essentially different kinds of Gaifman clique. Cliques of the first kind are those induced by an individual relational ground atom; any such clique is therefore guarded, and thus cannot be avoided in guarded bisimilar companion structures, not even in the guarded unravelling of the given structure.

Cliques of the second kind are incidental in the sense that several relational ground atoms play together to induce a clique configuration in the Gaifman graph, but no ground atom covers the entire clique. Such cliques are not guarded, they are obstacles for tree decompositions, and they would be broken up in the guarded unravelling of the given structure. Of course, the existence of this second kind of clique is precisely what makes CGF more expressive than plain GF.

When looking at hypergraphs rather than relational structures we mainly abstract away from the actual relational information and only retain the hypergraph structure induced by the guarded subsets, i.e., the subsets covered by single relational ground atoms. As remarked above, the hypergraph point of view directly relates to accessibility via guarded quantification, since the hyperedges are precisely the guarded subsets.

Note that in general this filtered view of a relational structure still retains strictly more information than the Gaifman graph, at least for vocabularies which have relations of arities greater than 2. One way to look at *conformality* of hypergraphs (in our case, the induced hypergraph), is that precisely in the conformal case the hypergraph structure is fully determined by the induced graph structure (in our case, the Gaifman graph). By definition, conformality means that the hyperedges are precisely the cliques (the guarded subsets are precisely the Gaifman cliques). In the world of guarded logics, then, a conformal hypergraph pattern corresponds to a structure in which clique guarded quantification is no stronger than ordinary guarded quantification, since all cliques are (covered by) hyperedges.

In this paper, we will construct finite conformal companions (covers) for finite hypergraphs. Our construction serves as a partial finite analogue of the generally infinite tree-like unravellings. As with unravellings, the relationship between these conformal covers and the base structure is simultaneously one of a "cover" and one of bisimilarity, mediated by local bijections governed by back-and-forth conditions with respect to hyperedges. The hypergraph construction lifts to the level of relational structures, where we obtain a cover which induces a guarded bisimulation, mediated by local isomorphisms governed by back-and-forth conditions with respect to guarded sets.

The central hypergraph construction is detailed in section 2, proving our main theorem for hypergraphs, Theorem 7. This section can be read independently, as it is based only on some preliminaries from sections 1.1, 1.2 and 1.3. The remainder of the paper links the main theorem to relational structures. The relational counterpart of the main theorem, in particular, is stated as Corollary 20 and further explored with respect to the two major applications in section 3.

As to the two applications mentioned above, section 3.3 deals with a reduction from the clique guarded fragment CGF to the guarded fragment GF, which (unlike guarded unravellings) is applicable in restriction to finite models. So we obtain a direct reduction for finite satisfiability and hence a new proof of the finite model property for CGF.

In section 3.1 we apply our results in the construction of finite extensions of partial automorphisms of the base structure to automorphisms of an extended structure (EPPA: extension property for partial automorphisms). Herwig's EPPA construction [11, 13], can be taken further to yield an extension whose only Gaifman cliques are the (unavoidable) automorphic images of Gaifman cliques already present in the base structure. See Theorem 9 for the statement of our main result in this context. Such a ramification of EPPA has repeatedly been pointed out as an interesting open problem, not least because of its bearing on the finite model property for CGF. Further corollaries provide a simplified route to the EPPA for the classes of K_n-free finite graphs and Henson digraphs, as well as the EPPA for the class of conformal finite relational structures of any relational type.

§1. Preliminaries.

1.1. Hypergraphs and relational structures. We consider hypergraphs $H = (A, S)$ where A is any set and $S \subseteq \mathcal{P}(A)$ any collection of subsets. The members of S are called hyperedges. Occasionally we write $S(H)$ for the set S of hyperedges in $H = (A, S)$.

Relational structures are denoted $\mathfrak{A} = (A, \bar{R})$, A the universe of $\mathfrak{A}$, $\bar{R}$ the tuple of relations as interpreted in $\mathfrak{A}$. We shall exclusively consider finite and purely relational vocabularies. The width of a relational vocabulary is the maximum of the arities of its relations.

We are mainly interested in the behaviour of relational structures with respect to guarded logics (see section 1.5), where only guarded subsets and tuples are directly accessible by quantification; the relevant basic definitions are given in the first of the two following definitions. Much of the guarded behaviour of relational structures is captured at the level of the associated hypergraph, which just describes the pattern of the accessible patches devoid of the actual relational information, as defined in the second definition below.

Definition 1. Let $\mathfrak{A} = (A, \bar{R})$ be a relational structure.

(i) A subset $s \subseteq A$ is *guarded* if s is a singleton or $s = \{ a : a \text{ in } \boldsymbol{a} \}$ for some $\boldsymbol{a} \in R$, R in $\bar{R}$.
(ii) A guarded subset is *maximally guarded* if it is not a proper subset of any other guarded subset.
(iii) A subset $s \subseteq A$ is *clique guarded* if for any two $a, a' \in s$ there is some guarded subset containing a and a'.
(iv) A k-tuple $\boldsymbol{a} \in A^k$ is (*clique*) *guarded* in $\mathfrak{A}$ if $\boldsymbol{a} \in s^k$ for some (clique) guarded set $s \subseteq A$.

Definition 2. The *hypergraph* associated with a relational structure $\mathfrak{A}$ is the hypergraph $H(\mathfrak{A}) = (A, S(\mathfrak{A}))$ where

$$S(\mathfrak{A}) = \{ s \subseteq A : s \text{ maximally guarded in } \mathfrak{A} \}.$$

As is common in hypergraph theory (cf. [3]), we associate an induced graph $G(H)$ with every hypergraph H. Note that if $H = H(\mathfrak{A})$ is the hypergraph induced by a relational structure $\mathfrak{A}$, then $G(H)$ is just the Gaifman graph $G(\mathfrak{A})$ associated with $\mathfrak{A}$ (cf. [5]).

DEFINITION 3. The *graph* associated with a hypergraph $H = (A, S)$ is the undirected graph $G(H) = (A, E)$ where

$$E = \{ (a, a') \in A^2 : a \neq a' \text{ and } a, a' \in s \text{ for some } s \in S \}.$$

DEFINITION 4.

(i) A hypergraph $H = (A, S)$ is called *conformal* if every clique of $G(H)$ is contained within some hyperedge of H.

(ii) A relational structure $\mathfrak{A}$ is called *conformal* if all its clique guarded subsets are contained in guarded sets, i.e., if $H(\mathfrak{A})$ is a conformal hypergraph.

1.2. Tree decompositions. The notion of conformality as expressed in Definition 4 comes from classical hypergraph theory, cf. [3]. The relational counterpart is just a direct analogue. Conformality is closely linked to the notion of tree-decomposability and acyclicity of hypergraphs. We briefly outline these connections for the sake of background and context, not because we shall directly draw on them in the sequel. A finite hypergraph is *tree-decomposable* if it is reducible to the empty hypergraph by repeated application of the following (cf. Graham's algorithm in [2] or the GYO-reducts in [6]):

- delete (from the set of hyperedges) some hyperedge that is contained within some other hyperedge;
- delete (from the universe and any hyperedge) some vertex a that is covered by at most one hyperedge.

An infinite hypergraph is tree-decomposable if all its finite induced sub-hypergraphs are.

It is not hard to see that conformality is necessary for tree-decomposability. Another necessary condition for H to be tree-decomposable is that its associated graph $G(H)$ (cf. Definition 3) is *chordal*: any cycle in $G(H)$ of length greater than 3 must have a chord, i.e., an edge linking two vertices that are not next neighbours along the cycle. Together these two conditions in fact characterise tree-decomposability: a hypergraph H is tree-decomposable if and only if it is conformal and (the associated graph $G(H)$ is) chordal, see, e.g., [2]. In the literature, conformal chordal hypergraphs are mostly called *acyclic*.

The concept of tree-decomposability of relational structures, which is of great importance in the theory of relational databases [2] and also figures prominently in the model theory of guarded logics [9], is closely related to the hypergraph theoretic notion. Indeed, the usual notion of tree-decomposability of relational structures can be captured as follows. A relational structure $\mathfrak{A}$ is *tree-decomposable* (of width k) if there is a tree-decomposable hypergraph (A, S) over the universe A of $\mathfrak{A}$ such that $S \supseteq S(\mathfrak{A})$ and $|s| \leqslant k + 1$ for all $s \in S$. This is equivalent to the characterisation that $\mathfrak{A}$ is tree-decomposable of width k if its Gaifman graph is tree-decomposable of width k, i.e., has tree width at most k, in the graph theoretic sense [4]. The notions of tree-decomposability and tree-width for relational structures are thus straightforward extensions of the underlying notions for plain graphs. Tree-decomposability of the hypergraph $H(\mathfrak{A})$ associated with a relational structure $\mathfrak{A}$, however, is a stronger notion. While arbitrary sets may be used as patches in a tree-decomposition of $\mathfrak{A}$ or of its Gaifman graph $G(\mathfrak{A})$, only hyperedges—i.e., guarded sets—are admissible as patches in tree-decompositions of $H(\mathfrak{A})$. Thus, tree-decomposability of $H(\mathfrak{A})$ implies tree-decomposability of $\mathfrak{A}$ (of tree width less than the width of the vocabulary of $\mathfrak{A}$), but not vice versa. Consider for instance a cycle of length n, with nodes $\{0, 1, \dots, n-1\}$ and edges $\{i, i+1\}$ (for $i < n-1$) and $\{n-1, 0\}$. This has tree width 2 as a graph (using patches $\{0, i, i+1\}$ $(0 < i < n-1)$ of size 3), but it is not tree-decomposable as a hypergraph.

1.3. Hypergraph bisimulations and covers. The following notion of bisimilarity between hypergraphs is the natural adaptation of the usual Ehrenfeucht–Fraïssé style notion of structural equivalence to the setting where "structure" is induced by hyperedges. Its relationship with the notion of guarded bisimilarity is apparent when we think of hypergraphs associated with relational structures, as will be made explicit in section 1.6, in particular Observation 16 and Lemma 19.

Definition 5. A *bisimulation* between hypergraphs $H = (A, S)$ and $H' = (A', S')$ is a non-empty collection Z of partial 1–1 maps between A and A' whose domains and ranges are hyperedges in H and H', respectively, with the following back-and-forth property with respect to hyperedges:

forth: if $p: s \to s'$ is in Z and if $t \in S$, then there is some $q: t \to t'$ in Z such that p and q agree on their common domain.

back: if $p: s \to s'$ is in Z and if $t' \in S'$, then there is some $q: t \to t'$ in Z such that p^{-1} and q^{-1} agree on their common domain.

We write $Z: H \sim H'$ if Z is a bisimulation between H and H', and just $H \sim H'$ if there is such a bisimulation.

Definition 6. A *cover* of a hypergraph H is a hypergraph $\hat{H}$ together with a surjective map $\pi: \hat{H} \to H$ which induces a hypergraph bisimulation:

$$Z(\pi): \hat{H} \sim H \quad \text{where } Z(\pi) = \{\, \pi\restriction\hat{s} : \hat{s} \in S(\hat{H}) \,\}.$$

We write $\pi: \hat{H} \sim H$ to indicate this.

It should be noted that the familiar (tree-like) unravelling of a hypergraph, which results in a conformal (indeed acyclic, i.e., conformal and chordal) hypergraph, actually yields an acyclic cover, which in general is infinite. At least as far as conformality is concerned, our main result provides a substitute that is applicable within the context of finite hypergraphs.

Theorem 7. *Every finite hypergraph H admits a cover $\pi: \hat{H} \sim H$ by a finite and conformal hypergraph $\hat{H}$.*

We do not know whether this can be strengthened to conformal and k-chordal covers for every k, where k-chordality would forbid chordless cycles of lengths less than k. Such a parameterised version of hypergraph acyclicity would be the most one can hope for in finite hypergraphs.

Before going into the details of our construction, which is sufficiently canonical to satisfy some additional nice automorphism properties, we prepare the stage for the two applications mentioned above—one dealing with the relationship between the guarded fragment GF and its more expressive extension CGF; the other one focusing on extension properties for partial automorphisms over finite relational structures.

Note. The following three sections, which pave the way for those two applications, can be read or skipped selectively without loss of coherence.

1.4. Extension properties for partial automorphisms. A partial isomorphism between relational structures $\mathfrak{A}$ and $\mathfrak{B}$ of the same type is a partial 1–1 map from A to B which induces an isomorphism between the substructures induced on its domain and range. We write $\mathrm{Part}(\mathfrak{A}, \mathfrak{B})$ for the set of partial isomorphisms between $\mathfrak{A}$ and $\mathfrak{B}$. For a single relational structure $\mathfrak{A}$, partial isomorphisms $p \in \mathrm{Part}(\mathfrak{A}, \mathfrak{A})$ are referred to as partial automorphisms. $\mathrm{Aut}(\mathfrak{A})$ stands for the automorphism group of $\mathfrak{A}$. If $\mathfrak{A}^* \supseteq \mathfrak{A}$ is an extension of $\mathfrak{A}$, we say that an automorphism $f \in \mathrm{Aut}(\mathfrak{A}^*)$ extends $p \in \mathrm{Part}(\mathfrak{A}, \mathfrak{A})$ if $p = f \restriction \mathrm{dom}(p)$.

For the second part of the following definition let # be some globally defined class of objects over relational structures of the type of $\mathfrak{A}$. Examples are (Gaifman-)edges, tuples in a specific relation R, guarded tuples, clique guarded tuples (Gaifman cliques). EPPA stands for Extension Property for Partial Isomorphisms.

DEFINITION 8. Let $\mathfrak{A} \subseteq \mathfrak{A}^*$.

(i) $\mathfrak{A}^*$ is an *EPPA extension* of $\mathfrak{A}$ if every $p \in \mathrm{Part}(\mathfrak{A}, \mathfrak{A})$ extends to an automorphism of $\mathfrak{A}^*$.

(ii) $\mathfrak{A}^*$ is a *#-faithful* extension of $\mathfrak{A}$ if every # over $\mathfrak{A}^*$ is the image of some # over $\mathfrak{A}$ under some automorphism of $\mathfrak{A}^*$.

(iii) A class of relational structures has the *#-faithful extension property for partial automorphisms*, or #-EPPA, if for every structure $\mathfrak{A}$ in that class there is a #-faithful EPPA extension $\mathfrak{A}^*$ also in that class.[1]

The fundamental EPPA results are the following. They are usually stated without particular attention to levels of faithfulness. However, one can prune an arbitrary EPPA extension in a straightforward way to obtain one that is faithful as stated. Indeed, if in an arbitrary EPPA extension $\mathfrak{A}^* \supseteq \mathfrak{A}$, we replace relation $R^{\mathfrak{A}^*}$ by the $\mathrm{Aut}(\mathfrak{A}^*)$-closure of $R^{\mathfrak{A}}$, then the resulting

[1]In fact this is a simplified version of EPPA, compared to the notion introduced in [13]; for the classes to be considered here, however, this causes no loss of generality.

structure is an EPPA extension of $\mathfrak{A}$ that is faithful with respect to tuples in R, cf. [13].

Hrushovski's EPPA Theorem [15]. The class of finite graphs has the EPPA; faithfulness with respect to edges is implicit. A greatly simplified and elegant proof of Hrushovski's theorem—combinatorial rather than group theoretic—is presented in [13, section 4.1].

Herwig's EPPA Theorem [11]. The class of finite relational structures (of any fixed finite relational type) has the EPPA; faithfulness with respect to guarded tuples (or guarded sets) is implicit.

We shall show the following by way of subjecting the result of a Herwig EPPA extension to a suitably adapted conformal cover construction.

THEOREM 9. *The class of finite relational structures* (*of any fixed finite relational type*) *has the Gaifman clique faithful EPPA*.

COROLLARY 10. *The following classes have the EPPA*:

(i) *Finite triangle-free graphs.* [11]
(ii) *Finite K_n-free graphs, for every $n \geqslant 3$.* [12]
(iii) *'Henson digraphs': finite directed graphs with no subgraph isomorphic to a tournament in $\mathcal{K}$, where $\mathcal{K}$ is an arbitrary class of finite tournaments.* [12]
(iv) *Finite conformal τ-structures, for every relational type τ.*

Note that (i)–(iii), even though not new, are here obtained by a much simpler and entirely combinatorial construction, if we use our construction on top of Lascar's simple construction for Hrushovski's theorem in [13, section 4.1]. Indeed, this new combinatorial approach to EPPA may be regarded as a further stepping stone in the methodological development of EPPA results so far, which has been markedly dual along group-theoretic versus combinatorial lines. While Hrushovski's original EPPA for graphs [15] as well as Herwig's generalisation to relational structures of higher arity [11] are group-theoretic, Lascar's proof of EPPA for graphs and its generalisation to higher arity in [13] are purely combinatorial. Now EPPA for K_n-free graphs and Henson digraphs, previously only available via group theory, gains an alternative, fully combinatorial account with the present construction.

OBSERVATION 11. *As our construction just adds a further extension layer, which moreover projects homomorphically onto its base structure, it also remains compatible with conditions concerning forbidden homomorphic images, in the sense of* [12] *and of the Herwig-Lascar ramification of Herwig's EPPA theorem*, [13].

1.5. Guarded logics. *Note*. This section and the next are not necessary for a coherent account of our hypergraph construction in section 2 and its EPPA application in section 3.1. Their topic will only be resumed in sections 3.2 and 3.3.

The guarded fragment of first-order logic, GF, was introduced by Andréka, van Benthem and Németi [1] as a first-order fragment capturing the spirit of modal quantification in the broader relational setting. Intuitively, in GF we may quantify over guarded tuples.

We write FO for first-order logic; for a first-order formula φ, $\mathrm{var}(\varphi)$ denotes the set of all variables occurring in φ, $\mathrm{free}(\varphi)$ the set of variables that have a free occurrence in φ.

DEFINITION 12. The formulae of GF $\subseteq$ FO are obtained inductively as the closure of atomic formulae (in a relational vocabulary τ, with equality) under Boolean connectives and the following quantification rules. For every $\varphi(\boldsymbol{x})$ in GF and every τ-atom $\alpha(\boldsymbol{x})$ and any tuple $\boldsymbol{y}$ such that $\{\, y : y \text{ in } \boldsymbol{y} \,\} \cup \mathrm{free}(\varphi) \subseteq \mathrm{var}(\alpha)$, the following are also formulae of GF:

$$\forall \boldsymbol{y}\ \big(\alpha(\boldsymbol{x}) \rightarrow \varphi(\boldsymbol{x})\big),$$
$$\exists \boldsymbol{y}\ \big(\alpha(\boldsymbol{x}) \wedge \varphi(\boldsymbol{x})\big).$$

The semantics is just the usual one for first-order logic.

Among the extensions of the guarded fragment that have since been considered we single out the following, which was introduced under the name of "clique guarded fragment", CGF, by Grädel in [8]. Under the name of "packed fragment", a similar extension of GF was considered by Marx [16]; indeed the packed and clique guarded fragments are syntactic variants of the same logic. Essentially, these fragments have quantification over clique guarded tuples rather than just guarded tuples. Correspondingly, CGF is known to be strictly more expressive than GF. In fact, CGF subsumes the loosely guarded fragment LGF of van Benthem [17]. LGF itself was there introduced as an important proper extension of GF to capture and generalise quantification patterns like the UNTIL construct in temporal logic. CGF further extends—arguably in the most natural way—the concept of loosely guarded quantification to the setting of vocabularies of width greater than 2.

The following syntactic conventions regarding CGF will for our purposes be superseded by the semantically equivalent ones to be given in Observation 14 below.

DEFINITION 13. The formulae of CGF $\subseteq$ FO are obtained inductively with the following more liberal quantification rule.

Let $\gamma(\boldsymbol{x})$ in free variables $\boldsymbol{x}$ be a conjunction $\bigwedge_{x,x' \in \boldsymbol{x}} \beta(x, x')$ over formulae $\beta(x, x')$ of the form $\exists \boldsymbol{z}\, \alpha(x, x', \boldsymbol{z})$, where α is an atom in which the displayed variables all occur, and the $\boldsymbol{z}$ is disjoint from the $\boldsymbol{x}$. (Semantically γ forces any instantiation of $\mathrm{free}(\gamma)$ to form a Gaifman clique.) Then, if $\varphi(\boldsymbol{x}) \in$ CGF and $\{\, y : y \text{ in } \boldsymbol{y} \,\} \cup \mathrm{free}(\varphi) \subseteq \mathrm{free}(\gamma)$, the following are also in CGF:

$$\forall \boldsymbol{y}\ \big(\gamma(\boldsymbol{x}) \rightarrow \varphi(\boldsymbol{x})\big),$$
$$\exists \boldsymbol{y}\ \big(\gamma(\boldsymbol{x}) \wedge \varphi(\boldsymbol{x})\big).$$

It is sometimes useful to resort to the following "normalisation" for guarded or clique guarded quantification. For a fixed finite vocabulary τ and variable tuple $\boldsymbol{x} = (x_1, \dots, x_n)$ we fix formulae $\mathrm{G}(\boldsymbol{x})$ and $\mathrm{CG}(\boldsymbol{x})$ which uniformly define the sets of those n-tuples that are guarded, respectively clique guarded in any τ-structure $\mathfrak{A}$:

$$\begin{aligned}\mathrm{G}[\mathfrak{A}] &= \{\, \boldsymbol{a} \in A^n : \mathfrak{A} \models \mathrm{G}[\boldsymbol{a}] \,\} &&= \{\, \boldsymbol{a} \in A^n : \boldsymbol{a} \text{ guarded in } \mathfrak{A} \,\},\\ \mathrm{CG}[\mathfrak{A}] &= \{\, \boldsymbol{a} \in A^n : \mathfrak{A} \models \mathrm{CG}[\boldsymbol{a}] \,\} &&= \{\, \boldsymbol{a} \in A^n : \boldsymbol{a} \text{ clique guarded in } \mathfrak{A} \,\}.\end{aligned}$$

Note that $\mathrm{G}(\boldsymbol{x})$ and $\mathrm{CG}(\boldsymbol{x})$ can actually both be formalised in GF, for instance,

$$\mathrm{CG}_0(\boldsymbol{x}) = \bigwedge\nolimits_{x,x' \in \boldsymbol{x}} \bigvee\nolimits_{\alpha} \exists \boldsymbol{z}\, \alpha(x, x', \boldsymbol{z}),$$

where α runs through all τ-atoms in which the displayed variables do occur and the $\boldsymbol{z}$ are from a fixed supply disjoint from $\boldsymbol{x}$. Some syntactic overhead can be avoided, however, in equivalent FO formalisations that do not adhere to the official syntax of GF, as for instance in

$$\mathrm{CG}(\boldsymbol{x}) = \bigwedge\nolimits_{x,x' \in \boldsymbol{x}} \bigvee\nolimits_{R} \exists \boldsymbol{z} \left(R\boldsymbol{z} \wedge \bigvee\nolimits_{z \in \boldsymbol{z}} x = z \wedge \bigvee\nolimits_{z \in \boldsymbol{z}} x' = z\right),$$

where R ranges over all relations in τ (and equality), and the $\boldsymbol{z}$ are as above. Similarly, for $\mathrm{G}(\boldsymbol{x})$ we may use

$$\mathrm{G}(\boldsymbol{x}) = \bigvee\nolimits_{R} \exists \boldsymbol{z} \left(R\boldsymbol{z} \wedge \bigwedge\nolimits_{x \in \boldsymbol{x}} \bigvee\nolimits_{z \in \boldsymbol{z}} x = z\right).$$

One can then characterise GF and CGF—or rather logics that are mere syntactic variants of them—through the stipulation of the following quantification rules. Translations between this modified syntax and the official standard are straightforward. The idea simply is to cover specific guards $\alpha(\boldsymbol{x})$ or $\gamma(\boldsymbol{x})$ for quantification in GF or CGF by the uniform guards $\mathrm{G}(\boldsymbol{x})$ or $\mathrm{CG}(\boldsymbol{x})$, as in replacing $\forall \boldsymbol{y} \left(\gamma(\boldsymbol{x}) \to \varphi(\boldsymbol{x})\right)$ by $\forall \boldsymbol{y} \left(\mathrm{CG}(\boldsymbol{x}) \to (\gamma(\boldsymbol{x}) \to \varphi(\boldsymbol{x}))\right)$.

Observation 14. *Let formulae* $\mathrm{G}(\boldsymbol{x})$ *and* $\mathrm{CG}(\boldsymbol{x})$ *globally define the sets of guarded, respectively clique guarded tuples* $\boldsymbol{x}$. *Then every formula of* GF *is logically equivalent to a first-order formula in which all quantifications are of the form indicated below, and vice versa.* (*We simultaneously introduce shorthand for relativised quantification.*)

$$\begin{aligned}(\forall \boldsymbol{y}.\mathrm{G}(\boldsymbol{x}))\, \varphi(\boldsymbol{x}) &:= \forall \boldsymbol{y} \left(\mathrm{G}(\boldsymbol{x}) \to \varphi(\boldsymbol{x})\right),\\ (\exists \boldsymbol{y}.\mathrm{G}(\boldsymbol{x}))\, \varphi(\boldsymbol{x}) &:= \exists \boldsymbol{y} \left(\mathrm{G}(\boldsymbol{x}) \wedge \varphi(\boldsymbol{x})\right),\end{aligned}$$

where $\{\, y : y \text{ in } \boldsymbol{y} \,\} \cup \mathrm{free}(\varphi) \subseteq \{\, x : x \text{ in } \boldsymbol{x} \,\}$.
Similarly for CGF *one uses the formulae* $\mathrm{CG}(\boldsymbol{x})$.

1.6. Guarded bisimulations. The Ehrenfeucht-Fraïssé equivalence associated to GF, guarded bisimulation equivalence, generalises bisimulation equivalence and is precisely adapted to capture quantification over guarded tuples. Also compare our hypergraph bisimulations in Definition 5.

DEFINITION 15. Let $\mathfrak{A}$ and $\mathfrak{A}'$ be relational structures of the same type. A *guarded bisimulation* between $\mathfrak{A}$ and $\mathfrak{A}'$ is a non-empty collection of partial isomorphisms $Z \subseteq \mathrm{Part}(\mathfrak{A}, \mathfrak{A}')$, where all $p \in Z$ have as domains/ranges guarded sets in $\mathfrak{A}/\mathfrak{A}'$, with the following back-and-forth property for guarded sets (s, t and s', t' guarded in $\mathfrak{A}$ and $\mathfrak{A}'$, respectively):

forth: for every $p\colon s \rightarrow s'$ in Z and every t there is some $q\colon t \rightarrow t'$ in Z such that p and q agree on their common domain.

back: for every $p\colon s \rightarrow s'$ in Z and every t' there is some $q\colon t \rightarrow t'$ in Z such that p^{-1} and q^{-1} agree on their common domain.

We write $Z\colon \mathfrak{A} \sim_{\mathrm{g}} \mathfrak{A}'$ if Z is a guarded bisimulation between $\mathfrak{A}$ and $\mathfrak{A}'$, and $\mathfrak{A} \sim_{\mathrm{g}} \mathfrak{A}'$ if there is such.

The similarity between hypergraph bisimulations and guarded bisimulations is apparent when we look at hypergraphs associated with relational structures.

OBSERVATION 16. *Any guarded bisimulation $Z\colon \mathfrak{A} \sim_{\mathrm{g}} \mathfrak{A}'$ induces a hypergraph bisimulation between the associated hypergraphs $Z\colon H(\mathfrak{A}) \sim H(\mathfrak{A}')$. Conversely, any hypergraph bisimulation $Z\colon H(\mathfrak{A}) \sim H(\mathfrak{A}')$ such that $Z \subseteq$* $\mathrm{Part}(\mathfrak{A}, \mathfrak{A}')$ *induces a guarded bisimulation. In other words: guarded bisimulations 'are' hypergraph bisimulations that respect the relational structure.*

The following semantic characterisation theorem is a central result from [1], underlining the role of guarded bisimulations and the naturalness of GF.

THEOREM 17. *For every first-order sentence φ in a relational vocabulary the following are equivalent*:

(i) *φ is invariant under guarded bisimulations*:

$$\mathfrak{A} \sim_{\mathrm{g}} \mathfrak{A}' \quad \textit{implies} \quad \mathfrak{A} \models \varphi \iff \mathfrak{A}' \models \varphi.$$

(ii) *φ is equivalent to a sentence of* GF.

DEFINITION 18. A *guarded cover* of a relational structure $\mathfrak{A}$ is a structure $\hat{\mathfrak{A}}$ of the same relational type as $\mathfrak{A}$ together with a surjective homomorphism $\pi\colon \hat{\mathfrak{A}} \rightarrow \mathfrak{A}$ which induces a guarded bisimulation:

$$Z(\pi)\colon \hat{\mathfrak{A}} \sim_{\mathrm{g}} \mathfrak{A} \quad \text{where } Z(\pi) = \{\, \pi\restriction\hat{s} : \hat{s} \in S(\hat{\mathfrak{A}}) \,\}.$$

We write $\pi\colon \hat{\mathfrak{A}} \sim_{\mathrm{g}} \mathfrak{A}$ to denote this.

It should be noted that the guarded unravelling of relational structures results in guarded covers by (generally infinite) structures $\hat{\mathfrak{A}}$ of bounded tree width, whose associated hypergraph $H(\hat{\mathfrak{A}})$ moreover is tree-decomposable. Grädel's *generalised tree model property* of guarded logic GF [9] actually applies in the stronger sense that every satisfiable sentence of GF has a model $\mathfrak{A}$ for which $H(\mathfrak{A})$ is tree-decomposable.

It is straightforward to see that guarded covers are related to bisimilar covers of hypergraphs in the following sense.

Lemma 19. *Let $\mathfrak{A}$ be a relational structure, $H = H(\mathfrak{A})$ the induced hypergraph. Then every hypergraph cover $\pi\colon \hat{H} \sim H$ induces a canonical guarded cover $\pi\colon \hat{\mathfrak{A}} \sim_{\mathrm{g}} \mathfrak{A}$, where $\hat{\mathfrak{A}}$ is a structure of the same relational type as $\mathfrak{A}$ over the universe $\hat{A}$ of $\hat{H}$, such that $H(\hat{\mathfrak{A}}) = \hat{H}$.*

One merely interprets all relations over the universe $\hat{A}$ of $\hat{H}$ so as to turn the restrictions $\pi{\restriction}\hat{s}$ into partial isomorphisms, for all hyperedges $\hat{s} \in \hat{S}$.

So we have the following corollary to Theorem 7.

Corollary 20. *Every finite relational structure $\mathfrak{A}$ admits a guarded cover $\pi\colon \hat{\mathfrak{A}} \sim_{\mathrm{g}} \mathfrak{A}$ where $\hat{\mathfrak{A}}$ is finite and conformal.*

§2. Construction of the conformal cover.

2.1. Conformal covers for hypergraphs. This section is devoted to the construction which proves Theorem 7. Fix a finite hypergraph $H = (A, S)$. We may assume that $A \notin S$, else H is already conformal. Let

$$U := \{\, u \subseteq A : \text{for all } s \in S,\ u \not\subseteq s \,\}.$$

So U is non-empty and consists precisely of the subsets of A that must not be the projections of cliques in the cover. Consider the product

$$A \times \prod_{u \in U} |u|,$$

where we identify a natural number n (in our case $n = |u|$, the cardinality of u) with the set $n = \{0, \dots, n-1\}$. A typical element of the product thus can be regarded as a pair (a, χ_a) whose second component is a function $\chi_a\colon U \to \mathbb{N}$ such that $\chi_a(u) < |u|$ for all $u \in U$. The universe $\hat{A}$ of the desired cover will consist of those elements (a, χ_a) for which

$$\chi_a(u) = 0 \iff a \notin u, \quad \text{for all } u \in U.$$

Note that this stipulation leaves a range $\{1, \dots, |u|-1\}$ of size smaller than $|u|$ for the values $\chi_a(u)$ whenever $a \in u$.

$$\hat{A} := \Big\{\, (a, \chi_a) \in A \times \prod_{u \in U} |u| : \text{for all } u \in U,\ \chi_a(u) = 0 \iff a \notin u \,\Big\}.$$

We let π be the natural projection $\pi\colon \hat{A} \to A$, given by $\pi(a, \chi_a) = a$.

We say that a subset $\hat{s} \subseteq \hat{A}$ is *generic* if

(i) $\pi{\restriction}\hat{s}$ is injective, i.e., for any two distinct elements (a, χ_a) and (b, χ_b) of $\hat{s}$ we have $a \neq b$.
(ii) for any two distinct elements (a, χ_a) and (b, χ_b) of $\hat{s}$ and $u \in U$: if a, $b \in u$ then $\chi_a(u) \neq \chi_b(u)$.

Note that (ii) can be rephrased as follows: for every u, the following function $\chi_{\_}(u)$ is injective (where (a, χ_a) is the element above a in $\hat{s}$):

$$\begin{aligned} \chi_{\_}(u)\colon u \cap \pi(\hat{s}) &\longrightarrow \{1, \dots, |u|-1\} \\ a &\longmapsto \chi_a(u). \end{aligned}$$

It follows immediately that the projection $\pi(\hat{s})$ of any generic set $\hat{s}$ must be contained in some hyperedge s of H. Otherwise we would have $\pi(\hat{s}) = u$ for some $u \in U$, and by condition (ii), for that u, $\chi_{\_}(u)\colon u \to \{1, \dots, |u| - 1\}$ would have to be an injection which is clearly absurd. For the hyperedges of $\hat{H}$ we now choose

$$\hat{S} := \big\{ \hat{s} \subseteq \hat{A} : \hat{s} \text{ generic and } \pi(\hat{s}) \in S \big\}.$$

In order to establish that $\pi\colon \hat{H} \sim H$ is indeed a cover, we claim that the system

$$Z = \big\{ \pi{\restriction}\hat{s} : \hat{s} \in \hat{S} \big\}\colon \hat{H} \sim H$$

is a hypergraph bisimulation. Note that the $\pi{\restriction}\hat{s}$ are bijections whose domains are hyperedges of $\hat{H}$ and whose ranges are hyperedges of H.

The forth-property is obvious. Consider the back-property for some $p = \pi{\restriction}\hat{s}\colon \hat{s} \to s = \pi(\hat{s})$, and some hyperedge $s' \in S$ of H. We need to find a generic set $\hat{s}'$ with $\pi(\hat{s}') = s'$ and such that $\pi(\hat{s} \cap \hat{s}') = s \cap s'$, or (equivalently) such that $\hat{s}' \cap \pi^{-1}(s \cap s') = \hat{s} \cap \pi^{-1}(s \cap s')$. Then $\pi{\restriction}\hat{s}'$ is as required by the back property. For the desired $\hat{s}'$ it remains to fix, for every $a \in s' \setminus s$, the values $\chi_a(u)$ such that $\chi_{\_}(u)\colon u \cap s' \to \{1, \dots, |u| - 1\}$ is injective for every single u. Consider a fixed $u \in U$. Observe that $u \not\subseteq s'$, whence $|u \cap s'| < |u|$. That part of $\chi_{\_}(u)$ over $u \cap s \cap s'$, which is already fixed, is injective by genericity of $\hat{s}$. We can therefore extend the injection $\chi_{\_}(u)$ from $u \cap s \cap s'$ to an injection over all of $u \cap s'$ as desired. This finishes the proof that $\pi\colon \hat{H} \sim H$ is a bisimilar cover.

Clearly $\hat{H}$ is conformal. For, if $\hat{s} \subseteq \hat{A}$ is a clique in $G(\hat{H})$, then any two distinct elements $(a, \chi_a) \neq (b, \chi_b)$ of $\hat{s}$ must be elements of some generic set in $\hat{A}$. It follows that $a \neq b$ and $\chi_a(u) \neq \chi_b(u)$ whenever $a, b \in u$. So $\hat{s}$ is a generic subset. Therefore $\pi(\hat{s}) \subseteq s'$ for some $s' \in S$. Arguing exactly as in the proof of the back-property for π, we find a hyperedge $\hat{s}'$ of $\hat{H}$ above s' for which $\hat{s} \subseteq \hat{s}'$. So every clique $\hat{s}$ is contained in a hyperedge of $\hat{H}$ as required.

This finishes the proof of the theorem. We collect some further useful properties of the cover $\pi\colon \hat{H} \sim H$ in the following lemma. These automorphism properties reflect on the canonical nature of our construction and will be essential for the EPPA application.

LEMMA 21. *Let the conformal cover $\pi\colon \hat{H} \sim H$ be obtained according to the above construction.*

(i) *For every $s \in S$ and generic sets $\hat{s}$ and $\hat{s}'$ above s (i.e., with $\pi(\hat{s}) = \pi(\hat{s}') = s$) there is an automorphism f of $\hat{H}$ that fixes all π-fibres setwise and maps $\hat{s}$ to $\hat{s}'$.*

(ii) *Every automorphism f of H admits a lift to an automorphism $\hat{f}$ of $\hat{H}$: $\hat{f} \in \mathrm{Aut}(\hat{H})$ and $\pi \circ \hat{f} = f \circ \pi$.*

For the first claim let $\hat{s}$, $\hat{s}'$, s be as stated. For every $a \in s$ let (a, χ_a) and (a, χ'_a) be the elements above a in $\hat{s}$ and $\hat{s}'$, respectively. For $u \in U$ consider the set

$$\{ (\chi_a(u), \chi'_a(u)) : a \in s \cap u \}.$$

Note that this set is the graph of a partial bijection ρ^0_u on $\{1, \dots, |u| - 1\}$. This is a direct consequence of the genericity of $\hat{s}$ and $\hat{s}'$: if $a, b \in u \cap s$ are distinct then so are $\chi_a(u)$ and $\chi_b(u)$ (genericity of $\hat{s}$) as well as $\chi'_a(u)$ and $\chi'_b(u)$ (genericity of $\hat{s}'$). Extend every ρ^0_u to a full permutation ρ_u of $|u|$ which fixes 0. It is easy to check that the following mapping, which is defined on all of $\hat{A}$, is a fibre-preserving automorphism that maps $\hat{s}$ to $\hat{s}'$:

$$(a, \chi_a) \mapsto (a, \chi'_a) \quad \text{where } \chi'_a(u) = \rho_u(\chi_a(u)).$$

For the second claim consider some $f \in \mathrm{Aut}(H)$. Then the following map $\hat{f}$ is as desired:

$$\hat{f}(a, \chi_a) := (b, \chi) \quad \text{where} \begin{cases} b = f(a) \\ \chi(u) = \chi_a(f^{-1}(u)). \end{cases}$$

That $\hat{f} \in \mathrm{Aut}(\hat{H})$ relies on the fact that f preserves S and hence also U.

2.2. Related covers for relational structures. While Corollary 20 directly follows from Theorem 7, we briefly indicate what one gets if instead of $H(\mathfrak{A})$ we use a different (in particular, richer) hypergraph over A to construct an induced cover.

Let $\mathfrak{A}$ be a relational structure and let $H = (A, S)$ be some hypergraph over the universe A of $\mathfrak{A}$. Let $\pi: \hat{H} \to H$ be a conformal cover, obtained as above. Interpret the relations in $\hat{\mathfrak{A}} = (\hat{A}, \dots)$ minimal under the requirement that the restrictions $\pi{\restriction}\hat{s}$ to hyperedges $\hat{s} \in S(\hat{H})$, and hence its restrictions to all generic subsets $\hat{s} \subseteq \hat{A}$, become partial isomorphisms:

$$R^{\hat{\mathfrak{A}}} := \bigcup_{\hat{s} \in S(\hat{H})} (\pi{\restriction}\hat{s})^{-1}\big(R^{\mathfrak{A}}{\restriction}\pi(\hat{s})\big).$$

LEMMA 22. *For $\pi: \hat{\mathfrak{A}} \to \mathfrak{A}$ as above*:

(i) $\pi: \hat{\mathfrak{A}} \to \mathfrak{A}$ *is a surjective homomorphism; moreover, π is a partial isomorphism in restriction to every generic subset $\hat{s} \subseteq \hat{A}$, and the system of the $\pi{\restriction}\hat{s}$ has the back-and-forth properties with respect to the hyperedges of H and the generic subsets of $\hat{A}$.*

(ii) *For every $s \in S$, $\mathfrak{A}{\restriction}s$ embeds isomorphically into $\hat{\mathfrak{A}}$, with generic image in $\hat{A}$.*

(iii) *For all $s \in S$ and generic sets $\hat{s}$ and $\hat{s}'$ with $\pi(\hat{s}) = \pi(\hat{s}') = s$ there is an automorphism of $\hat{\mathfrak{A}}$ that fixes all π-fibres set-wise and maps $\hat{s}$ to $\hat{s}'$.*

(iv) *Every automorphism of $\mathfrak{A}$ that preserves S admits a lift to an automorphism of $\hat{\mathfrak{A}}$. In particular, if S is invariant under* $\mathrm{Aut}(\mathfrak{A})$, *then every automorphism of $\mathfrak{A}$ lifts to $\hat{\mathfrak{A}}$.*

(v) *If $H = (A, S)$ is such that $S \supseteq S(\mathfrak{A})$, then $\pi: \hat{\mathfrak{A}} \to \mathfrak{A}$ is a guarded cover.*

The first two claims are obvious from the construction. Claims (iii) and (iv) are strictly analogous to the corresponding claims in Lemma 21 above. Claim (v) uses the fact that every guarded subset of $\hat{\mathfrak{A}}$ is by construction contained in a generic subset of $\hat{A}$; $S \supseteq S(\mathfrak{A})$ implies that the guarded subsets of $\mathfrak{A}$ are contained in hyperedges of H; the claim then follows by the last observation in (i).

§3. Two applications.

3.1. Clique faithful EPPA. We prove Theorem 9. A Gaifman clique faithful EPPA extension is obtained from a Herwig EPPA as a suitable relational cover as follows.

Given $\mathfrak{A}$, let $\mathfrak{B} \supseteq \mathfrak{A}$ be a Herwig EPPA extension:

(i) Every partial automorphism p of $\mathfrak{A}$ extends to an automorphism of $\mathfrak{B}$.
(ii) Every guarded tuple (or guarded set, or tuple in a relation) of $\mathfrak{B}$ is the image under some automorphism of $\mathfrak{B}$ of a guarded tuple (or guarded set, or tuple in a relation) in $\mathfrak{A}$.

Now apply the relational cover construction of section 2.2 based on the hypergraph $H = (B, S)$ where

$$S = \{ s \subseteq B : f(s) \subseteq A \text{ for some } f \in \mathrm{Aut}(\mathfrak{B}) \}.$$

Let $\pi\colon \hat{\mathfrak{B}} \to \mathfrak{B}$ be the resulting cover. As $A \subseteq B$ forms a hyperedge of H, $\mathfrak{A}$ embeds isomorphically into $\hat{\mathfrak{B}}$, according to Lemma 22 (ii). Let $\rho\colon \mathfrak{A} \to \hat{\mathfrak{B}}$ be this embedding, $\hat{\mathfrak{A}} = \rho(\mathfrak{A}) \subseteq \hat{\mathfrak{B}}$ the isomorphic image of $\mathfrak{A}$ whose universe $\hat{A} \subseteq \hat{B}$ is a generic subset. We further claim that now

(iii) every Gaifman clique of $\hat{\mathfrak{B}}$ is $\mathrm{Aut}(\hat{\mathfrak{B}})$-related to some Gaifman clique of $\hat{\mathfrak{A}}$.
(iv) every partial automorphism of $\hat{\mathfrak{A}}$ extends to an automorphism of $\hat{\mathfrak{B}}$.

These claims prove that $\hat{\mathfrak{B}}$ provides a Gaifman clique faithful EPPA extension of $\mathfrak{A}$, if, without loss of generality we identify $\mathfrak{A}$ with its isomorphic image $\hat{\mathfrak{A}}$ in $\hat{\mathfrak{B}}$.

For (iii), it follows from the general construction that every Gaifman clique $\hat{s}$ in $\hat{\mathfrak{B}}$ forms a generic set in $\hat{B}$, whence its projection is contained within some $s \in S$. By the choice of S, there is an automorphism f of $\mathfrak{B}$ which maps s into $A \subseteq B$. By Lemma 22 (iv) (our set S of hyperedges is invariant under $\mathrm{Aut}(\mathfrak{B})$), there is a lift $\hat{f}$ of f to an automorphism of $\hat{\mathfrak{B}}$. $\hat{f}$ being a lift we find that $\pi(\hat{f}(\hat{s})) \subseteq A$. Therefore $\hat{f}(\hat{s})$ and $\rho(f(s)) \subseteq \hat{A}$ are two generic sets above $f(s) \subseteq A$. By Lemma 22 (iii) there is an automorphism g of $\hat{\mathfrak{B}}$ which maps $\hat{f}(\hat{s})$ to $\rho(f(s))$. The composition $g \circ \hat{f}$ therefore maps $\hat{s}$ into $\hat{A}$. Clearly then, $g(\hat{f}(\hat{s}))$ is a Gaifman clique in $\hat{\mathfrak{A}}$.

(iv) is proved in a similar way. Let $\hat{p} \in \mathrm{Part}(\hat{\mathfrak{A}}, \hat{\mathfrak{A}})$. Clearly $\hat{p}$ is the lift and isomorphic image of a partial isomorphism p of $\mathfrak{A}$. $\mathfrak{B}$ has an automorphism f which extends p. Let $\hat{f}$ be a lift of f to $\hat{\mathfrak{B}}$, obtained according to Lemma

22 (iv). Choose g as a fibre-preserving automorphism that maps $\hat{f}(\mathrm{dom}(\hat{p}))$ to range$(\hat{p})$ (both generic above range$(p) \subseteq A$), according to Lemma 22 (iii). Then $g \circ \hat{f}$ extends $\hat{p}$.

The same argument yields the ramified version of Observation 11, if correspondingly we start with an EPPA extension $\mathfrak{B} \supseteq \mathfrak{A}$ which does omit homomorphic images of some finite collection $\mathcal{K}$ of structures. This additional requirement is trivially preserved, since $\pi \colon \hat{\mathfrak{B}} \to \mathfrak{B}$ is a homomorphism, cf. Lemma 22 (i).

3.2. FMP and EPPA: the general pattern. With the clique faithful variant of EPPA one can prove the finite model property (FMP) for CGF in complete analogy with the proof given by Grädel [9] for GF, which is based on the use of Herwig's EPPA. Indeed, the issue whether or not CGF also enjoys the FMP had been one motivation to look for a clique faithful strengthening of Herwig's construction. This latter issue remained open even when the FMP for CGF was settled in [14]. Before we proceed (in the next section) to give a new and direct reduction from FINSAT(CGF) to FINSAT(GF), it may be useful to survey the general pattern in which (Gaifman clique or guarded tuple) faithful EPPA extensions give rise to finite models for (clique or ordinarily) guarded formulae.

We treat GF and CGF in parallel to highlight the uniform nature of the argument. Let φ be in GF or CGF, syntactically presented according to Observation 14 (normalised quantification pattern). In both cases we firstly use a Skolemisation procedure with respect to (guarded or clique guarded) subformulae of φ. In effect this means that we merely have to deal with formulae of the form:

$$\big(\forall \boldsymbol{z}.\#(\boldsymbol{z})\big)\ \big(\beta(\boldsymbol{z}) \wedge \big(\exists \boldsymbol{y}.\#(\boldsymbol{x})\big)\ \gamma(\boldsymbol{x})\big) \quad \text{where} \begin{cases} \#(\boldsymbol{x}) = \mathrm{CG}(\boldsymbol{x}) & \text{for CGF} \\ \#(\boldsymbol{x}) = \mathrm{G}(\boldsymbol{x}) & \text{for GF} \end{cases}$$

where β, γ are quantifier-free, $\mathrm{var}(\gamma) \subseteq \{\, x : x \text{ in } \boldsymbol{x}\,\}$, $\mathrm{var}(\beta) \subseteq \{\, z : z \text{ in } \boldsymbol{z}\,\}$, $\boldsymbol{y}$ and $\boldsymbol{z}$ disjoint, $\boldsymbol{x} = \boldsymbol{y}\boldsymbol{z}$ (we do not rule out the case that either $\boldsymbol{z}$ or $\boldsymbol{y}$ is empty, though).

Note. Skolemisation for clique guarded subformulae of a given φ introduces a new Skolem predicate for each such subformula. These new predicates must not serve as guards, though. In other words the formulae $\mathrm{CG}(\boldsymbol{x})$ are fixed and only pertain to clique guardedness in (the reducts to) the original vocabulary of φ.

Starting from a (supposedly infinite) model $\mathfrak{B}$ of the Skolemisation of φ we obtain a finite model as a #-faithful EPPA extension of a sufficiently rich finite substructure of $\mathfrak{B}$, as follows. Let $\mathfrak{A} \subseteq \mathfrak{B}$ be finite and such that all isomorphism types of (small) #-substructures of $\mathfrak{B}$ are represented as substructures of $\mathfrak{A}$. In the case of CGF, more precisely, we limit the size of the clique guarded substructures under consideration to the maximal width of clique guards $\mathrm{CG}(\boldsymbol{x})$ in φ. This is important as Gaifman cliques (unlike guarded sets) can be of unbounded size even for fixed finite vocabulary.

We then get, for any #-faithful EPPA extension $\mathfrak{A}^* \supseteq \mathfrak{A}$:

$$\mathfrak{B} \models \big(\forall \boldsymbol{z}.\#(\boldsymbol{z})\big)\ \big(\beta(\boldsymbol{z}) \wedge \big(\exists \boldsymbol{y}.\#(\boldsymbol{x})\big)\ \gamma(\boldsymbol{x})\big)$$
$$\iff \mathfrak{A}^* \models \big(\forall \boldsymbol{z}.\#(\boldsymbol{z})\big)\ \big(\beta(\boldsymbol{z}) \wedge \big(\exists \boldsymbol{y}.\#(\boldsymbol{x})\big)\ \gamma(\boldsymbol{x})\big).$$

For this claim, consider an instantiation $\boldsymbol{a}$ for $\boldsymbol{z}$ in $\mathfrak{A}^*$ with $\mathfrak{A}^* \models \#(\boldsymbol{a})$. As $\mathfrak{A}^*$ is #-faithful over $\mathfrak{A}$ we may without loss of generality assume that $\boldsymbol{a}$ is in $\mathfrak{A}$.

Therefore, a counter example to $\big(\forall \boldsymbol{z}.\#(\boldsymbol{z})\big)\ \beta(\boldsymbol{z})$ in $\mathfrak{A}^*$ would immediately give a counter example in $\mathfrak{A}$ and hence in $\mathfrak{B}$.

Similarly, if $\mathfrak{B} \models \big(\forall \boldsymbol{z}.\#(\boldsymbol{z})\big)\ \big(\exists \boldsymbol{y}.\#(\boldsymbol{x})\big)\ \gamma(\boldsymbol{x})$ and we consider $\boldsymbol{a}$ in $\mathfrak{A}$ such that $\mathfrak{A}^* \models \#(\boldsymbol{a})$ we know that $\mathfrak{B}, \boldsymbol{a} \models \big(\exists \boldsymbol{y}.\#(\boldsymbol{x})\big)\ \gamma(\boldsymbol{x})$. Let $\boldsymbol{b}$ be such that $\mathfrak{B}, \boldsymbol{a}, \boldsymbol{b} \models \#(\boldsymbol{x}) \wedge \gamma(\boldsymbol{x})$. Let $\mathfrak{B}_0 = \mathfrak{B}{\restriction}\boldsymbol{ab}$. By the choice of $\mathfrak{A}$ we have an isomorphic copy $(\mathfrak{B}_0', \boldsymbol{a}', \boldsymbol{b}')$ of $(\mathfrak{B}_0, \boldsymbol{a}, \boldsymbol{b})$ inside $\mathfrak{A}$. Let p be the partial automorphism of $\mathfrak{A}$ that maps $\boldsymbol{a}'$ to $\boldsymbol{a}$. Let f be an extension of p to an automorphism of $\mathfrak{A}^*$. Then $\mathfrak{A}^*, \boldsymbol{a}, f(\boldsymbol{b}') \models \#(\boldsymbol{x}) \wedge \gamma(\boldsymbol{x})$, and hence $\mathfrak{A}^*, \boldsymbol{a} \models \big(\exists \boldsymbol{y}.\#(\boldsymbol{x})\big)\ \gamma(\boldsymbol{x})$ as desired.

3.3. A reduction from FINSAT(CGF) to FINSAT(GF). Finally, we present the promised translation from CGF to GF, based on Skolemisation for Gaifman cliques. This translation $\varphi \mapsto \varphi^*$ will be such that every model of φ can be expanded to a model of φ^*. Moreover, φ^* implies φ over conformal structures. This translation then obviously serves as a reduction from SAT(CGF) to SAT(GF), because we merely have to unravel a model of φ^* to obtain a conformal model of φ^*, which then also is a model of φ. Thus, we obtain:

$$\varphi \in \mathrm{SAT(CGF)} \iff \varphi^* \in \mathrm{SAT(GF)}.$$

In more detail, from $\varphi \in$ CGF we obtain $\varphi^* \in$ GF as follows. For formulae of CGF and GF we appeal to the syntactic normalisation of Observation 14. Let r be the maximal width of clique guards $\mathrm{CG}(\boldsymbol{x})$ occurring in φ and let R_c be a new relation symbol of arity r. Let $\mathrm{G}^*(\boldsymbol{x})$ be formulae expressing guardedness with respect to the extended vocabulary, R_c inclusive.

In φ we replace any clique guarded quantifications of the form

$$(\forall \boldsymbol{y}.\mathrm{CG}(\boldsymbol{x}))\ \xi(\boldsymbol{x}) \quad \text{or} \quad (\exists \boldsymbol{y}.\mathrm{CG}(\boldsymbol{x}))\ \xi(\boldsymbol{x})$$

by their simply guarded forms over the extended vocabulary

$$(\forall \boldsymbol{y}.\mathrm{G}^*(\boldsymbol{x}))\ \xi(\boldsymbol{x}) \quad \text{or} \quad (\exists \boldsymbol{y}.\mathrm{G}^*(\boldsymbol{x}))\ \xi(\boldsymbol{x}),$$

respectively. Finally add, as a conjunct in φ^*, a formalisation in GF of

$$\varphi_0 := \forall \boldsymbol{x}\ \big(R_c\boldsymbol{x} \rightarrow \mathrm{CG}(\boldsymbol{x})\big),$$

where $\mathrm{CG}(\boldsymbol{x})$ is for the old vocabulary, R_c exclusive.

Clearly every model of φ has an expansion that is a model of φ_0 and hence of φ^*: interpret R_c as the set of all clique guarded tuples of arity r.

Conversely, a model of φ^* is also a model of φ provided it does not have "false cliques". If $(\mathfrak{A}, R_c) \models \varphi^*$, then all tuples in R_c are indeed Gaifman cliques (i.e., clique guarded) in $\mathfrak{A}$ as $(\mathfrak{A}, R_c) \models \varphi_0$. But $\mathfrak{A}$ could have "false" Gaifman cliques that are not covered by any guarded set of $(\mathfrak{A}, R_c)$. If, however, $(\mathfrak{A}, R_c)$ is conformal then $(\mathfrak{A}, R_c) \models \varphi^*$ implies $\mathfrak{A} \models \varphi$. Indeed, in conformal models of φ_0, all clique guarded tuples are guarded, whence $\mathrm{G}^*(\boldsymbol{x})$, $\mathrm{CG}(\boldsymbol{x})$, $\mathrm{CG}^*(\boldsymbol{x})$ are all equivalent: $\mathrm{G}^*(\boldsymbol{x}) \Longrightarrow \mathrm{CG}(\boldsymbol{x})$ uses φ_0; $\mathrm{CG}(\boldsymbol{x}) \Longrightarrow \mathrm{CG}^*(\boldsymbol{x})$ is trivial; $\mathrm{CG}^*(\boldsymbol{x}) \Longrightarrow \mathrm{G}^*(\boldsymbol{x})$ is conformality.

Therefore, as the usual $\sim_{\mathrm{g}}$-unravelling of any structure yields a guarded bisimilar structure that is conformal (albeit generally infinite), we see that φ is satisfiable if and only if φ^* is satisfiable. As Corollary 20 similarly provides finite conformal guarded bisimilar companion structures, we find that in complete analogy also $\varphi \in \mathrm{FINSAT}(\mathrm{CGF}) \iff \varphi^* \in \mathrm{FINSAT}(\mathrm{GF})$.

COROLLARY 23. *The translation* $\varphi \mapsto \varphi^*$ *provides a simultaneous reduction from* $\mathrm{SAT}(\mathrm{CGF})$ *to* $\mathrm{SAT}(\mathrm{GF})$ *and from* $\mathrm{FINSAT}(\mathrm{CGF})$ *to* $\mathrm{FINSAT}(\mathrm{GF})$.

In particular, the finite model property for GF, [9], directly implies the finite model property for CGF, [14].

Note. The above translation is polynomial if we use a succinct formalisation of $\mathrm{G}(\boldsymbol{x})$ as indicated in connection with Observation 14. Passage through non-standard syntax for GF and CGF may be avoided with an alternative translation, which can also be kept polynomial in terms of the official syntax for both GF and CGF. Syntactic normalisation according to Observation 14 has been chosen here for the sake of clarity and to enable a more uniform presentation of the central idea in the translation. These considerations may be important if one wants to make sure that the translation is compatible with the complexities established for $\mathrm{SAT}(\mathrm{CGF})$ and $\mathrm{SAT}(\mathrm{GF})$ in [8, 9].

REFERENCES

[1] HAJNAL ANDRÉKA, JOHAN VAN BENTHEM, and ISTVAN NÉMETI, *Modal languages and bounded fragments of predicate logic*, ***Journal of Philosophical Logic***, vol. 27 (1998), pp. 217–274.

[2] CATRIEL BEERI, RONALD FAGIN, DAVID MAIER, and MIHALIS YANNAKAKIS, *On the desirability of acyclic database schemes*, ***Journal of the Association for Computing Machinery***, vol. 30 (1983), pp. 479–513.

[3] CLAUDE BERGE, ***Graphs and hypergraphs***, North-Holland, 1973.

[4] REINHARD DIESTEL, ***Graph theory***, 2nd ed., Springer-Verlag, 2000.

[5] HEINZ-DIETER EBBINGHAUS and JÖRG FLUM, ***Finite model theory***, Springer-Verlag, 1995.

[6] GEORG GOTTLOB, NICOLA LEONE, and FRANCESCO SCARCELLO, *The complexity of acyclic conjunctive queries*, ***Journal of the Association for Computing Machinery***, vol. 43 (2001), pp. 431–498.

[7] ERICH GRÄDEL, *The decidability of guarded fixed point logic*, ***JFAK. Essays dedicated to Johan van Benthem on the occasion of his 50th birthday*** (J. Gerbrandy, M. Marx, M. de Rijke, and Y. Venema, editors), CD-ROM, Amsterdam University Press, 1999.

[8] ———, *Decision procedures for guarded logics*, ***Proceedings of 16th international conference on automated deduction, CADE 16, 1999***, Lecture Notes in Computer Science, vol. 1632, Springer-Verlag, 1999, pp. 31–51.

[9] ———, *On the restraining power of guards*, ***The Journal of Symbolic Logic***, vol. 64 (1999), pp. 1719–1742.

[10] ERICH GRÄDEL, COLIN HIRSCH, and MARTIN OTTO, *Back and forth between guarded and modal logics*, ***Association for Computing Machinery Transactions on Computational Logic***, vol. 3 (2002), pp. 418–463.

[11] BERNHARD HERWIG, *Extending partial isomorphisms on finite structures*, ***Combinatorica***, vol. 15 (1995), pp. 365–371.

[12] ———, *Extending partial isomorphisms for the small index property of many (omega)-categorical structures*, ***Israel Journal of Mathematics***, vol. 107 (1998), pp. 93–124.

[13] BERNHARD HERWIG and DANIEL LASCAR, *Extending partial isomorphisms and the profinite topology on free groups*, ***Transactions of the American Mathematical Society***, vol. 352 (2000), pp. 1985–2021.

[14] IAN HODKINSON, *Loosely guarded fragment of first-order logic has the finite model property*, ***Studia Logica***, vol. 70 (2002), pp. 205–240.

[15] EHUD HRUSHOVSKI, *Extending partial isomorphisms of graphs*, ***Combinatorica***, vol. 12 (1992), pp. 411–416.

[16] MAARTEN MARX, *Tolerance logic*, ***Journal of Logic, Language and Information***, vol. 10 (2001), pp. 353–374.

[17] JOHAN VAN BENTHEM, *Dynamic bits and pieces*, ***Technical Report LP-97-01***, ILLC, Universiteit van Amsterdam, 1997.

DEPARTMENT OF COMPUTING
IMPERIAL COLLEGE LONDON
LONDON SW7 2AZ, UK
E-mail: imh@doc.ic.ac.uk
URL: www.doc.ic.ac.uk/~imh

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF WALES
SWANSEA, SA2 8PP, UK
E-mail: m.otto@swan.ac.uk
URL: www-compsci.swan.ac.uk/~csmartin