

Bisimulation and Coverings for Graphs and Hypergraphs

Martin Otto

Department of Mathematics
Technische Universität Darmstadt

Abstract. We survey notions of bisimulation and of bisimilar coverings both in the world of graph-like structures (Kripke structures, transition systems) and in the world of hypergraph-like general relational structures. The provision of finite analogues for full infinite tree-like unfoldings, in particular, raises interesting combinatorial challenges and is the key to a number of interesting model-theoretic applications.

1 Introduction

Bisimulation provides the fundamental back&forth methodology to capture dynamic notions of behavioural equivalence in terms structural similarity. On one hand, bisimulation stands in the tradition of model-theoretic comparison games and may be viewed as a specialisation and adaptation of the classical Ehrenfeucht–Fraïssé method to a specifically modal context in the analysis of transition systems and graph-like structures, cf. [10]. On the other hand, bisimulation captures the quintessential notion of game equivalence, and can be seen as a general framework for the wider family of back&forth games that pertain to various other types of structures with their specific notions of observable configurations (game positions) and accessibility (transitions/moves) [21]. With respect to such generalisations, a very natural and seemingly basic, yet also very powerful, generalisation step takes us

- from graph-like structures (vertices linked by edges) to hypergraphs (subsets linked by their overlaps/intersections);
- from ordinary bisimulation to hypergraph or guarded bisimulation;
- from modal logics to guarded logics.

This generalisation, and the interesting parallels as well as the combinatorial and model-theoretic challenges it entails are at the heart of this presentation.

Both for graphs and for hypergraphs, bisimulation equivalence supports a natural process of tree unfolding [8, 10, 12]:

- ordinary tree unfoldings of transition systems or graphs into tree structures of finite edge-labelled paths from a designated root vertex;
- tree-like hypergraph unfoldings of hypergraphs (or relational structures) into acyclic hypergraphs (or relational structures).

These unfoldings are linked to the original structures not just by some bisimulation, but they provide bisimilar *coverings*,¹ which project homomorphically to the original structure and allow for lifts of all available links to related loci in the covering structure. In other words, the bisimulation that links the covering to the base structure is induced by the homomorphic projection of the cover: the homomorphism condition guarantees the *forth*-property for the bisimulation relationship, while the *back*-property corresponds to the lifting requirement. The resulting unfoldings are acyclic in the ordinary sense of graph theory (no cycles, not even undirected ones) or of hypergraph theory (α -acyclic, the strongest notion of hypergraph acyclicity).

The mere existence of acyclic bisimilar unfoldings has strong model-theoretic and algorithmic consequences for related logics – in fact for any logic whose semantics is preserved in the process of unfolding. The *tree model property* [8] for modal logics and the *generalised tree model property* [11] of guarded logics account for various decidability and complexity results for modal and guarded logics; they stand behind the applicability of tree automata to their algorithmic model theory; they support model-theoretic characterisations of modal and guarded fragments of first-order and second-order logics; and, even though the unfoldings themselves are typically infinite, they also point us in the right direction for understanding certain approaches to the *finite model property* and the combinatorial challenges posed by the finite model theory of modal and guarded logics [10, 19, 26, 23, 11, 13].

2 Bisimulations and Coverings

At the combinatorial level, we speak of graphs and hypergraphs. A hypergraph is a structure of the format $\mathcal{A} = (A, S)$ with vertex set A and set of hyperedges $S \subseteq \mathcal{P}(A)$; the case of graphs is the special case in which all hyperedges have cardinality ≤ 2 (so these are undirected, but not necessarily loop-free, graphs). The graphs and hypergraphs of interest arise as combinatorial abstractions of relational structures:

¹ I prefer this classical term to the more idiosyncratic terminology of ‘bounded morphisms’ in the modal logic tradition [8].

- transition systems or Kripke structures (relational structures of width ≤ 2) as graph-like structures, in which vertices and edges carry labels (edge- and vertex-coloured graphs) and edges may be directed; the graph abstraction is the underlying undirected link pattern between vertices, including loops.
- relational structures $\mathcal{A} = (A, (R^A))$ with relations R of arbitrary, fixed arities; the hypergraph abstraction is the hypergraph of guarded subsets, which has, for each tuple $\mathbf{a} = (a_1, \dots, a_r) \in R^A$, the set of its components $[\mathbf{a}] := \{a_1, \dots, a_r\}$ as one of its hyperedges, and in addition, every singleton set. It is often convenient to close the set of hyperedges under downward inclusion.

We define bisimulations and coverings for graphs and hypergraphs. They naturally arise in the study of relational structures, in which case the homomorphism needs to respect relational content: colour of vertices and colours and direction of edges in the case of graph-like structures;² and the full isomorphism type of the induced substructures on hyperedges in the general case. We leave this to the imagination of the reader and just note that coverings of the underlying graph or hypergraph structures of relational structures can always be lifted to proper coverings of those relational structures in a canonical manner. For graph-like structures one may also wish to disentangle multiple edges (edges between the same vertices that are of different colours). The traditional modal variant of bisimulation equivalence for transition systems of Kripke structures is, moreover, directional; and for this, one would formulate the back&forth requirements, in a setting of directed graph-like structures, as a lifting condition for forward extensions only.

Definition 2.1. [graph covering] A *covering* of a graph $\mathcal{A} = (A, E)$ by a graph $\hat{\mathcal{A}} = (\hat{A}, \hat{E})$ is a homomorphism $\pi: \hat{\mathcal{A}} \rightarrow \mathcal{A}$ such that, for every $\hat{a} \in \hat{A}$ and edge $e = (a, a') \in E$ incident at $a = \pi(\hat{a})$, there is some edge $\hat{e} \in \hat{E}$ with $\pi(\hat{e}) = e$ (we regard \hat{e} as a lifting of e to \hat{a}). The covering is *faithful* if it preserves incidence degrees in the sense that there always is a unique lifting of e to \hat{a} .

Definition 2.2. [hypergraph covering] A *covering* of a hypergraph $\mathcal{A} = (A, S)$ by a hypergraph $\hat{\mathcal{A}} = (\hat{A}, \hat{S})$ is a map (hypergraph homomorphism) $\pi: \hat{\mathcal{A}} \rightarrow \mathcal{A}$ such that the restrictions $\pi \upharpoonright \hat{s}$ of π to the hyperedges $\hat{s} \in \hat{S}$

² Note that a loop in a transition system is bisimilar to, and may be covered by, a non-trivial cycle or an infinite chain; in this sense, the graph case is not quite the immediate specialisation of the hypergraph case.

are bijections onto hyperedges $\pi(\hat{s}) = s \in S$, and such that, for every $\hat{s} \in \hat{S}$ above $s = \pi(\hat{s})$ and every $s' \in S$, there is some hyperedge $\hat{s}' \in \hat{S}$ with $\pi(\hat{s}') = s'$ and $\pi \upharpoonright (\hat{s} \cap \hat{s}') = s \cap s'$.

Note that the *back*-condition in hypergraph coverings says that the overlap between s and s' can be lifted to any \hat{s} above s .

The following notions of acyclicity are the natural ones; for hypergraphs we concentrate on the strongest of the common notions of acyclicity, also known as α -acyclity [6, 5]. The formulation below, in terms of conformality and chordality, is one of several equivalent ones [5]: in particular, it is equivalent to the existence of a tree decomposition whose bags are the hyperedges.

It is convenient to associate with a hypergraph $\mathcal{A} = (A, S)$ an induced graph $G(\mathcal{A})$ on the same vertex set A in which distinct vertices a and a' are linked by an edge if they are elements of the same hyperedge $s \in S$, i.e., $G(\mathcal{A})$ is a superposition of cliques replacing the hyperedges of \mathcal{A} .

Definition 2.3. [acyclicity] A graph is *acyclic* if it does not have any cycles; a hypergraph $\mathcal{A} = (A, S)$ is *acyclic* if it is both

- (i) *conformal*: each clique in $G(\mathcal{A})$ is contained in a single hyperedge, and
- (ii) *chordal*: every cycle in $G(\mathcal{A})$ of length greater than 3 has a chord, i.e., $G(\mathcal{A})$ has no induced subgraphs isomorphic to the k -cycle for $k > 3$.

Every graph $\mathcal{A} = (A, E)$ possesses natural coverings by acyclic graphs \mathcal{A}_a^* , whose nodes are the paths from some designated root vertex a ; these coverings can be made faithful by restriction to non-degenerate edge-labelled paths, i.e., to edge-labelled paths that do not reverse onto themselves. It is immediate that all acyclic coverings of a graph that is itself not acyclic must be infinite.

It is not hard to see that, similarly, every hypergraph $\mathcal{A} = (A, S)$ admits coverings by acyclic hypergraphs. Such coverings can be based on acyclic graph unfoldings of the associated intersection graph $I(\mathcal{A}) = (S, \Delta)$ where $\Delta = \{(s, s') : s \neq s', s \cap s' \neq \emptyset\}$. Again, acyclic coverings of a hypergraph that is itself not acyclic must be infinite.

We note that non-trivial coverings of hypergraphs can in general not preserve incidence degrees, i.e., they are usually branched coverings and there is no immediate analogy to faithful coverings, which are generally available only for graphs.

2.1 Model-theoretic applications

The availability of (maybe necessarily infinite) acyclic unfoldings, implies special model properties for logics whose semantics is invariant under

corresponding notions of bisimulation. Many facets of model-theoretic and algorithmic well-behavedness of modal and guarded logics can be attributed more or less directly to their characteristic invariances under bisimulation equivalence. We here concentrate on special model properties and issues of expressive completeness. For background on the logics involved, modal logics (including many variants and extensions of basic modal logic ML, e.g., in the spirit of epistemic or temporal logics) and guarded logics like the guarded fragment of first-order logic GF, we refer the reader to textbook sources and surveys [8, 10] and the original papers [2, 11, 14, 12].

Theorem 2.1. *All bisimulation-invariant logics of graph-like structures, and in particular basic modal logic ML and the modal μ -calculus L_μ , have the tree model property: every satisfiable formula is satisfiable in a tree-like model.*

All guarded bisimulation-invariant logics of relational structures, and in particular the guarded fragment GF and its fixpoint extension μ GF, have the generalised tree model property: every satisfiable formula is satisfiable in a model whose hypergraph of guarded subsets is acyclic.

Suitable tree-like models and their analysis, and the theory of tree automata that can be employed for model checking purposes, account also for characterisation theorems that characterise L_μ and μ GF as expressively complete for bisimulation-invariant properties that can be expressed in monadic second-order logic and guarded second-order logic respectively. For the following see the fundamental paper by Janin and Walukiewicz [17] and, for a reduction to [17] that extends the result to the hypergraph setting and guarded logics, [12].

Theorem 2.2. *A monadic second-order definable property of pointed transition systems is expressible in the modal μ -calculus if, and only if, it is invariant under modal bisimulation equivalence.*

A guarded second-order definable property of relational structures is expressible in guarded fixpoint logic if, and only if, it is invariant under guarded bisimulation equivalence.

The classical variants of these characterisations for the bisimulation-invariant fragments of first-order logic can be proved classically, by means of compactness arguments and saturated models. These results are due to van Benthem [28] and Andr eka, van Benthem and N emeti [2], respectively.

Theorem 2.3. *A first-order definable property of pointed transition systems is expressible in basic modal logic ML if, and only if, it is invariant under modal bisimulation equivalence.*

A first-order definable property of relational structures is expressible in the guarded fragment GF if, and only if, it is invariant under guarded bisimulation equivalence.

Returning to special model properties, it is interesting that basic modal logic ML, many of its first-order variants (and even the modal μ -calculus, but only as long as only forward modalities are considered), and also the guarded fragment (but not guarded fixpoint logic) have the *finite model property*, see textbook sources [8, 10] for ML, and [11] for GF. Finite tree models for formulae of ML can be obtained by pruning from tree-unfoldings of arbitrary models; finite models for formulae of GF can be obtained as natural combinatorial completions (w.r.t. extension properties for partial isomorphisms based on Herwig’s Theorem [15]) from truncations of (tree-unfoldings of) arbitrary models according to [11].

3 Acyclicity versus Finiteness

As outlined above, acyclic coverings of finite structures are necessarily infinite if the given finite structure is not itself acyclic – both in the world of graphs and of hypergraphs. However, graphs and hypergraphs admit finite coverings that are acyclic as far as size-bounded induced sub-configurations in the covering structure are concerned.

Definition 3.1. [N -acyclicity] For $N \in \mathbb{N}$, a graph or hypergraph \mathcal{A} is *N -acyclic* if every induced substructure of \mathcal{A} of size up to N is acyclic.

The (induced) sub-hypergraph $\mathcal{A} \upharpoonright B$ of $\mathcal{A} = (A, S)$, for some $B \subseteq A$, is the hypergraph structure obtained by restricting both the vertex set and each individual hyperedge to B : $\mathcal{A} \upharpoonright B = (B, S \upharpoonright B)$ where $S \upharpoonright B = \{s \cap B : s \in S\}$.

A graph is N -acyclic if its *girth* is greater than N . Also note that an N -acyclic graph for $N \geq 2\ell + 1$ is *ℓ -locally acyclic* in the sense that the induced subgraph on the ℓ -neighbourhood of any vertex is acyclic. This does not hold for hypergraphs: consider the hypergraph formed by the subdivision of a triangle with central vertex into the three triangles formed by the central vertex together with one edge each, of the outer triangle; it is clear that any non-trivial covering of this simple hypergraph will have a vertex above the central one whose 1-neighbourhood contains a chordless cycle of some length $3k$ or is infinite.

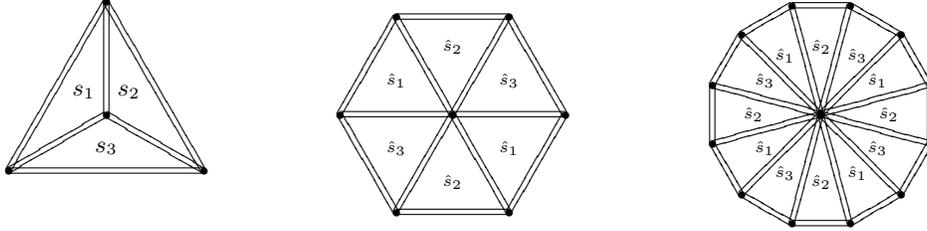


Fig. 1. Local unfoldings of a sub-divided triangle, 2-fold and 4-fold.

The discrepancy between locality in the associated graph $G(\mathcal{A})$ and locality in the intersection graph $I(\mathcal{A})$ is a characteristic feature of hypergraphs that account for their much more challenging combinatorics w.r.t. degrees of acyclicity in finite covers.

3.1 Cayley Groups and Finite Graph Coverings

N -acyclic graph covers can easily be obtained as direct products with highly regular, generic N -acyclic graphs that arise as Cayley graphs of suitable groups. Cayley graphs are a well-known source for regular graphs of large girth [1, 7].

If G is a finite group generated by the finite set of non-trivial involutive group elements $\{e: e \in E\} \subseteq \{g \in G: g^2 = 1 \neq g\} \subseteq G$, then the *Cayley graph* of G w.r.t. this generator set E is the undirected loop-free E -edge-coloured graph formed by the group elements with an edge of colour $e \in E$ precisely between all pairs of group elements of the form $(g, g \cdot e)$.

Conversely, let $\mathcal{A} = (A, (R_e)_{e \in E})$ be any finite E -edge-coloured undirected graph in which each R_e is a partial matching. We here call such graphs E -graphs; if we add an e -coloured loop at every vertex not incident with an e -coloured edge, then we obtain a *complete E -graph* with the property that every vertex is incident with precisely one edge (or loop) of colour e , for every $e \in E$.

A finite complete E -graph $\mathcal{A} = (A, (R_e))$ induces a finite Cayley group that arises as a subgroup $\text{sym}(\mathcal{A}) \subseteq \text{Sym}(A)$ of the symmetric group $\text{Sym}(A)$ on the set A . For this, each edge colour $e \in E$ is associated with a generator $e \in \text{sym}(\mathcal{A})$ according to

$$e: A \longrightarrow A \\ a \longmapsto a' \quad \text{if } (a, a') \in R_e,$$

which is a non-trivial involution provided \mathcal{A} has at least one e -coloured edge that is not a loop.

It is not hard to see that the girth of the Cayley graph of this group $\text{sym}(\mathcal{A})$ w.r.t. the generators $e \in E$ is at least $4k + 2$ if \mathcal{A} is chosen to be (the completion of) the full E -coloured tree of depth k .

To obtain an N -acyclic finite covering of an arbitrary finite graph $\mathcal{A} = (A, E)$, let G be a Cayley group of girth greater than N with involutive generators $e \in E$ (the edge set of \mathcal{A} , where we now think of each undirected edge, or loop, as the *set* of incident vertices, so that $e^{-1} = e$). We obtain a faithful N -acyclic covering by the natural direct product of \mathcal{A} and (the Cayley graph of) G :

$$\begin{aligned} \mathcal{A} \otimes G &:= (A \times G, E \otimes G) \\ \text{where } E \otimes G &= \{((a, g), (a', g \cdot e)) : \{a, a'\} = e \in E\}. \end{aligned}$$

This direct product provides a covering with the natural projection onto the first factor as the covering homomorphism. For details of the following we refer the reader to [18, 23]. Several specialisations to special classes of graph-like structures are discussed in [9].

Proposition 3.1. *For every $N \in \mathbb{N}$, every finite graph-like structure \mathcal{A} admits faithful coverings by finite N -acyclic structures.*

Cayley groups of much more than just large girth in the above sense are constructed in [23]. An iteration of amalgamation steps that produce more and more highly acyclic E -graphs and passages from these E -graphs \mathcal{A} to induced Cayley groups $\text{sym}(\mathcal{A})$ produces Cayley groups without *coset cycles* of length up to N (rather than just no generator cycles of length up to N , see Definition 3.2 below).

In a Cayley group with generator set E , and for subsets $\alpha \subseteq E$, we write $G[\alpha] \subseteq G$ for the subgroup generated by $\alpha \subseteq E$, and, for a group element $g \in G$, $gG[\alpha] = \{g \cdot h : h \in G[\alpha]\}$ for its coset w.r.t. this subgroup. In terms of the Cayley graph, $gG[\alpha]$ consists of those vertices that are reachable from g on α -coloured paths.

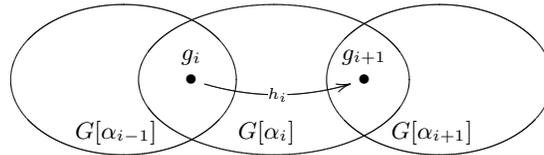


Fig. 2. Coset cycles: local view.

Definition 3.2. [coset cycles] Let G be a Cayley group with involutive generators $e \in E$. A *coset cycle* of length n in G consists of a tuple $(g_i)_{i \in \mathbb{Z}_n}$ in G and a tuple of subsets $(\alpha_i)_{i \in \mathbb{Z}_n}$ of E (cyclically indexed in \mathbb{Z}_n) such that for all i :

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

G is called *N -acyclic* if it does not have any coset cycles of length up to N .

Proposition 3.2. *For every $N \in \mathbb{N}$ and every finite set E there are finite N -acyclic Cayley groups with involutive generator set E .*

For the proof see [23], where such Cayley groups are used as auxiliary ingredients in the construction of N -acyclic hypergraph covers as we shall discuss in the next section.

3.2 Finite Hypergraph Coverings

An intricate iterative construction in [23], which involves finite truncated pieces of just locally finite N -acyclic coverings and their completion through gluing operations that do not create new short cycles, yields the following. In that construction, reduced products with N -acyclic Cayley groups (as discussed further below, see Lemma 3.1) are used as part of the completion steps needed to repair deficiencies along the boundaries where necessary extensions got cut off in the truncation.

Proposition 3.3. *For every $N \in \mathbb{N}$, every finite hypergraph (or relational structure) admits coverings by finite N -acyclic hypergraphs (relational structures whose hypergraph of guarded subsets is N -acyclic).*

The role of the N -acyclic Cayley graphs from Proposition 3.2 in this construction lies in the following preservation property in reduced products – which is quite unlike the role of Cayley groups of large girth towards the construction of N -acyclic graph covers.

Let $\mathcal{A} = (A, S)$ be a hypergraph and $\sigma: E \rightarrow \mathcal{P}(A)$ a finite family of regions $\sigma(e)$ to be used as glueing sites in A . If G is a Cayley group with involutive generators $e \in E$, we define the *reduced product* of \mathcal{A} with G (over σ) to be the hypergraph $\mathcal{A} \otimes_\sigma G = (\hat{A}, \hat{S})$, as follows.

The vertex set \hat{A} of $\mathcal{A} \otimes_\sigma G$ is obtained from the direct product $A \times G$ by identifications w.r.t. the equivalence relation \approx induced by

$$(a, g) \approx (a', g \cdot e) \quad \text{if} \quad a = a' \in \sigma(e).$$

The set of hyperedges \hat{S} of $\mathcal{A} \otimes_{\sigma} G$ consists simply of those subsets that are represented by subsets of the form $s \times \{g\}$ in the direct product, for $s \in S$. The natural projection π onto the first factor is well defined and turns $\pi: \hat{A} \rightarrow \mathcal{A}$ into a hypergraph cover. This construction and the following lemma are part of the rather involved route to N -acyclic hypergraph coverings in [23] that support Proposition 3.3.

Lemma 3.1. *If $\mathcal{A} = (A, S)$ is an N -acyclic hypergraph and if the finite set E parametrises glueing sites $\sigma(e) \subseteq A$ that are guarded by hyperedges of \mathcal{A} in the sense that $\sigma(e) \subseteq s$ for some $s \in S$, then the reduced product $\mathcal{A} \otimes_{\sigma} G$ with any N -acyclic Cayley group G with involutive generator set E is again N -acyclic.*

A much more transparent construction of finite hypergraph coverings, which also achieves feasible size bounds, is obtained in [3], where the following weaker degrees of acyclicity are explored.

Definition 3.3. [weak N -acyclicity] Let $N \in \mathbb{N}$ and consider two hypergraphs $\hat{\mathcal{A}} = (\hat{A}, \hat{S})$ and $\mathcal{A} = (A, S)$ in a hypergraph covering $\pi: \hat{A} \rightarrow \mathcal{A}$. This covering is *weakly N -acyclic* if every induced sub-hypergraph of $\hat{\mathcal{A}}$ of size up to N can be augmented by extra hyperedges which project onto singletons or subsets of hyperedges of \mathcal{A} so as to become acyclic: for $\hat{A}_0 \subseteq \hat{A}$ with $|\hat{A}_0| \leq N$ there is some $\hat{S}' \subseteq \{\hat{s}' \subseteq \hat{A}_0: \pi(\hat{s}') \subseteq s \in S \text{ or } |\pi(\hat{s}')| = 1\}$ such that $(\hat{A}_0, \hat{S} \upharpoonright \hat{A}_0 \cup \hat{S}')$ is acyclic.

The width $w(\mathcal{A})$ of a hypergraph $\mathcal{A} = (A, S)$ is defined to be the maximal size of hyperedges in S . The following core result from [3] was inspired by previous constructions in [25] as well as earlier and cruder approximations in [20]. Essentially, N -acyclic coverings are realised in [3] by finite quotients of suitable term structures that are designed to satisfy exactly the *back*-requirements imposed by the overlap pattern of the desired hypergraph.

Proposition 3.4. *For every $N \in \mathbb{N}$, every finite hypergraph \mathcal{A} admits weakly N -acyclic coverings by finite hypergraphs whose size can be bounded polynomially in $|\mathcal{A}|$ for fixed N and width of \mathcal{A} .*

In fact, for fixed width, the construction in [3] also produces coverings by conformal hypergraphs of size polynomial in the given hypergraph, thus improving on a simpler but less succinct construction in [16].

3.3 Model-Theoretic Applications

The following strengthenings of the finite model property for the guarded fragment GF are almost immediate as corollaries to Propositions 3.3 and 3.4, respectively. Details as well as interesting ramifications are presented in [23, 3].

In the following, a class of structures defined in terms of finitely many *forbidden finite substructures* is specified by a finite set of finite structures, and consists of all structures in which the specified structures do not occur as induced substructures, up to isomorphism. A class of structures defined in terms of *forbidden homomorphisms* from finitely many finite structures is similarly defined.

Corollary 3.1. *GF has the finite model property in restriction to any class of relational structures that is defined in terms of finitely many forbidden finite cyclic substructures.*

Corollary 3.2. *GF has the finite model property in restriction to any class of relational structures that is defined in terms of forbidden homomorphisms from finitely many finite structures.*

Among the many other consequences of the core construction of weakly N -acyclic coverings in [3] are improved bounds on the sizes of small finite models for sentences of the guarded fragment GF and its generalisation to the clique guarded fragment, and a further application that shows that finite relational structures admit polynomial time canonisation w.r.t. guarded bisimulation equivalence.

The primary model-theoretic goal in [23], on the other hand, is the analogue of the expressive completeness result in Theorem 2.3 for the guarded fragment, in the sense of finite model theory. Just as for basic modal logic ML, the assertion of expressive completeness in restriction to finite model theory does not follow (in any meaningful direct way, that is) from the classical fact, and its proof must follow very different, more constructive lines.

The finite N -acyclic coverings of [23] can be saturated w.r.t. branching behaviour so as to produce finite structures for which guarded bisimulation equivalence to sufficiently high finite depth entails first-order equivalence up to some given quantifier rank. See [23, 22] for details and the general context. This yields the following.

Theorem 3.1. *A first-order definable property of finite (!) relational structures is expressible in GF if, and only if, it is invariant among finite structures (!) under guarded bisimulation equivalence.*

4 Outlook

Among the latest developments in the directions outlined above are the following.

Restricted back&forth homomorphism equivalences: Back&forth equivalences that combine modal or guarded bisimulation (of graph-like or hypergraph-like relational structures, respectively) with local homomorphisms (of bounded size) under the name of *unary negation bisimulation* [27] and *guarded negation bisimulation* [4], have been investigated as the underlying semantic invariances of yet more expressive well-behaved fragments of first-order logic:

Unary and guarded negation fragments: The *unary negation fragment* UN and *guarded negation fragment* GN of first-order logic are based on unconstrained (or just size-constrained) positive existential quantification and the limitation of negation to formulae that have just a single free variable (unary negation) or an explicitly guarded tuple of free variables (guarded negation). These fragments essentially extend modal logic and the guarded fragment, respectively, and display very similar model-theoretic and algorithmic properties, see [27] and [4]. Technically, the generalisation of good model-theoretic features, and especially of the finite model property, involves reductions to Corollary 3.2 above. An analogue of Theorem 3.1 has recently been outlined in [22].

Improved hypergraph constructions: New, as yet unpublished constructions based on groupoids and their analogue of Cayley graphs allow for far more generic and direct constructions of suitably acyclic hypergraps and hypergraph coverings by means of reduced products [24].

References

1. N. Alon. Tools from higher algebra. In R. Graham, M. Grötschel, and Lovász, editors, *Handbook of Combinatorics*, volume II, pages 1749–1783. North-Holland, 1995.
2. H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.
3. V. Bárány, G. Gottlob, and M. Otto. Querying the guarded fragment. *Logical Methods in Computer Science*, 2012. submitted.
4. V. Bárány, B. ten Cate, and L. Segoufin. Guarded negation. In *Proceedings of ICALP*, pages 356–367, 2011.
5. C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30:497–513, 1983.

6. C. Berge. *Graphs and Hypergraphs*. North-Holland, 1973.
7. N. Biggs. Cubic graphs with large girth. In G. Blum et al., editor, *Annals of the New York Academy of Sciences*, volume 555, pages 56–62. 1989.
8. P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
9. A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *Annals of Pure and Applied Logic*, 161:1–42, 2009.
10. V. Goranko and M. Otto. Model theory of modal logic. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, pages 249–329. Elsevier, 2007.
11. E. Grädel. On the restraining power of guards. *Journal of Symbolic Logic*, 64:1719–1742, 1999.
12. E. Grädel, C. Hirsch, and M. Otto. Back and forth between guarded and modal logics. *ACM Transactions on Computational Logics*, 3:418 – 463, 2002.
13. E. Grädel and M. Otto. The freedoms of (guarded) bisimulation. in preparation, 2012.
14. E. Grädel and I. Walukiewicz. Guarded fixed point logic. In *Proc. LICS'99*, pages 45–54, 1999.
15. B. Herwig. Extending partial isomorphisms on finite structures. *Combinatorica*, 15:365–371, 1995.
16. I. Hodkinson and M. Otto. Finite conformal hypergraph covers and Gaifman cliques in finite structures. *Bulletin of Symbolic Logic*, 9:387–405, 2003.
17. D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ -calculus with respect to monadic second-order logic. In *Proceedings of CONCUR*, pages 263–277, 1996.
18. M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004.
19. M. Otto. Bisimulation invariance and finite models. In *Colloquium Logicum 2002*, Lecture Notes in Logic, pages 276–298. ASL, 2006.
20. M. Otto. Avoiding incidental homomorphisms into guarded covers. Technical report no. 2600, TU Darmstadt, 2009.
21. M. Otto. Model theoretic methods for fragments of FO and special classes of (finite) structures. In C. Steinhorn J. Esparza, C. Michaux, editor, *Finite and Algorithmic Model Theory*, volume 379 of *LMS Lecture Notes*, pages 271–341. CUP, 2011.
22. M. Otto. Expressive completeness through logically tractable models. *Annals of Pure and Applied Logic*, 2012. under revision.
23. M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59 (1), 2012.
24. M. Otto. On groupoids and hypergraphs. unpublished draft, 2012.
25. R. Rosati. On the decidability and finite controllability of query processing in databases with incomplete information. In *Proc. PODS'06*, pages 356–365, 2006.
26. E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.
27. B. ten Cate and L. Segoufin. Unary negation. In *Proceedings of STACS*, pages 344–355, 2011.
28. J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Napoli, 1983.