Avoiding Incidental Homomorphisms Into Guarded Covers

Martin Otto*
Technische Universität Darmstadt

November 16, 2009

Abstract

For a given finite relational structure, we want to construct a finite guarded-bisimilar companion structure (a guarded cover) that avoids all avoidable homomorphic images of other structures. 'Avoidable' turns out to mean that the given homomorphism does not lift to guarded unravellings of the original structure, hence can at least be eliminated in infinite guarded-bisimilar companion structures. We thus get the finite model property for an extension of the guarded fragment by 'forbidden homomorphic images.'

1 Preliminaries

Let τ be a finite, purely relational vocabulary of width k. For any τ-structure $\mathfrak{A}$ let

$$S[\mathfrak{A}] = \big\{ s \subseteq A \colon s \text{ maximally guarded in } \mathfrak{A} \big\}$$

be the set of guarded subsets of $\mathfrak{A}$ that are not properly contained in another guarded subset. Recall that $s \subseteq A$ is guarded in $\mathfrak{A}$ if it is a singleton or $s = \{a \colon a \in \mathbf{a}\}$ for some $\mathbf{a} \in R^{\mathfrak{A}}$ for some $R \in \tau$. We also consider

$$S{\downarrow}[\mathfrak{A}] = \big\{ s' \subseteq A \colon s' \subseteq s \text{ for some } s \in S[\mathfrak{A}] \big\}.$$

A hypergraph consists of a universe together with a collection of non-empty subsets of the universe. We refer to

$$\mathfrak{H}(\mathfrak{A}) = \big(A, S[\mathfrak{A}]\big)$$

as the *(guarded) hypergraph* associated with $\mathfrak{A}$. Beside this hypergraph we consider also two induced graphs:

$$\begin{aligned} \mathfrak{G}(\mathfrak{A}) &= \big(A, E\big) & \quad E &= \big\{ (a, a') \colon a \neq a', \{a, a'\} \in S{\downarrow}[\mathfrak{A}] \big\} \\ \mathfrak{S}(\mathfrak{A}) &= \big(S[\mathfrak{A}], E\big) & \quad E &= \big\{ (s, s') \colon s \neq s', s \cap s' \neq \emptyset \big\}. \end{aligned}$$

Classical terminology in hypergraph theory (see Berge) has the following.

- $\mathfrak{H}(\mathfrak{A})$ is *conformal* if for any $s_1, s_2, s_3 \in S[\mathfrak{A}]$ there is some $s \in S[\mathfrak{A}]$ such that $s \supseteq \bigcup_{i \neq j} s_i \cap s_j$. Equivalently: every clique of $\mathfrak{G}(\mathfrak{A})$ is contained in some $s \in s[\mathfrak{A}]$.

---

*This note is based on results first presented at an FMT meeting in Cambridge 2001, revised in 2008.

- $\mathfrak{H}(\mathfrak{A})$ is *chordal* if every simple cycle of length $l > 3$ in $\mathfrak{G}(\mathfrak{A})$, $a_0, \ldots, a_l$ distinct where $(a_i, a_{i+1}) \in E$ for all i (cyclic indexing), has a chord: $(a_i, a_j) \in E$ for some i, j with $|i - j| \not\equiv 1$.

By extension, we also say that $\mathfrak{G}(\mathfrak{A})$ or just $\mathfrak{A}$ are conformal or chordal.

2 Unravellings and covers

Recall the notion of a tree decomposition.

Definition 1 A tree $T = (V, E)$ with a mapping $\lambda \colon V \to \mathcal{P}(A)$ provides a tree decomposition of $\mathfrak{A}$ if

(i) every $s \in S[\mathfrak{A}]$ is covered by some $\lambda(v)$.

(ii) the local connectivity condition is satisfied: for all $a \in A$ the set $\{v \in V : a \in \lambda(v)\}$ is connected in T.

A tree decomposition T, λ is here called a *guarded tree decomposition* if $\lambda(v) \in S[\mathfrak{A}]$.

We refer to the sets $\lambda(v) \subseteq A$ as the *patches* of the tree decomposition (also commonly called *bags* of the decomposition).

Note that, for instance, any cycle has a tree decomposition with size 3 patches, but, since size 3 patches are necessary, no proper cycle of length at least 3 in $\mathfrak{G}(\mathfrak{A})$ can have a guarded tree decomposition in a structure $\mathfrak{A}$ with only unary and binary relations (width 2 vocabulary).

Definition 2 (i) A *cover* of $\mathfrak{A}$ is a structure $\mathfrak{B}$ which is guarded bisimilar to $\mathfrak{A}$ via a guarded bisimulation induced by a homomorphism $\pi \colon \mathfrak{B} \to \mathfrak{A}$. I.e., $\{\pi \restriction s \colon s \in S[\mathfrak{B}]\}$ is a system of partial isomorphisms satisfying the usual back-and-forth conditions. We write $\pi \colon \mathfrak{B} \sim_G \mathfrak{A}$ to indicate this special kind of guarded bisimulation.

(ii) An *unravelling* of $\mathfrak{A}$ is a guarded tree decomposable cover, i.e., a structure $\mathfrak{A}^*$ which admits a guarded tree decomposition and is guarded bisimilar to $\mathfrak{A}$ via a guarded bisimulation induced by a homomorphism $\pi \colon \mathfrak{A}^* \sim_G \mathfrak{A}$.[1]

Unravellings are important as they provide guarded tree decomposable companions up to guarded bisimulation. Every $\mathfrak{A}$ has an unravelling.

A canonical unravelling may be based on the tree $T = (V, E)$ of all sequences in $S[\mathfrak{A}]$, rooted at the empty sequence $\emptyset$. With a sequence $v = s_0 \ldots s_l$ associate the structure $\mathfrak{A}_v = \mathfrak{A} \restriction s_l$; form the disjoint union of the $\mathfrak{A}_v$ for all $v \in V$, and take for $\mathfrak{A}^*$ its quotient w.r.t. the congruence relation induced by the transitive closure of the identity relation on overlapping patches $\mathfrak{A}_v$ and $\mathfrak{A}_w$ where w is a child of u. Formally, $a \in \mathfrak{A}_u$ and $a \in \mathfrak{A}_v$ are identified iff u and v are in the same connected component of $\mathfrak{F}$ and $a \in \mathfrak{A}_u$ for all u that occur along a (or: the unique shortest) path from v to w in T. Let a_v stand for the equivalence class (under this congruence) of the element $a \in \mathfrak{A}_v$; let $\pi \colon a_v \mapsto a$ be the natural projection. The maximally guarded sets in $\mathfrak{A}^*$ are the sets $[v] = \big\{[a_v] \colon a \in \mathfrak{A}_v\big\}$. $\pi \restriction [v] \colon \mathfrak{A}^* \restriction [v] \simeq \mathfrak{A}_v \subseteq \mathfrak{A}$ is a partial isomorphism. It is not hard to check that, moreover, the back-and-forth property is satisfied so that $\pi \colon \mathfrak{A}^* \sim_G \mathfrak{A}$.

[1] Note that we here speak of *an* unravelling rather than the (canonical) unravelling (see below).

Lemma 3 *Any unravelling of* $\mathfrak{A}$ *can be homomorphically mapped into any* $\mathfrak{B} \sim_G \mathfrak{A}$*: if* $\pi\colon \mathfrak{A}^* \sim_G \mathfrak{A}$ *and* $\mathfrak{B} \sim_G \mathfrak{A}$*, then there is homomorphism* $h\colon \mathfrak{A}^* \to \mathfrak{B}$*.*

In fact, the claim of the lemma follows from the following more fundamental claim. If $\mathfrak{A} \sim_G \mathfrak{B}$ and $\mathfrak{A}$ is guarded tree decomposable, then there is a homomorphism $h\colon \mathfrak{A} \to \mathfrak{B}$. Even if $\mathfrak{A}$ itself is not guarded tree decomposable, an unravelling $\mathfrak{A}^*$ of $\mathfrak{A}$ is, and therefore the statement of the lemma follows.

To prove the claim, let $T = (V, E)$, $\lambda\colon V \to S[\mathfrak{A}]$ provide a guarded tree decomposition of $\mathfrak{A}$. Let $Z\colon \mathfrak{A} \sim_G \mathfrak{B}$ where $Z \subseteq \big\{p \in \mathrm{Part}(\mathfrak{A}, \mathfrak{B})\colon \mathrm{dom}(p) \in S[\mathfrak{A}], \mathrm{im}(p) \in S[\mathfrak{B}]\big\}$. We find a mapping $\lambda^*\colon V \to Z$ such that $\mathrm{dom}(\lambda^*(v)) = \lambda(v) \subseteq \mathfrak{A}$ and such that $\lambda^*(u)$ and $\lambda^*(v)$ agree on the intersection of their domains whenever $(u, v) \in E$. It is then easy to see that we may obtain the desired homomorphism $h\colon \mathfrak{A} \to \mathfrak{B}$ by putting $h(a) = \big(\lambda^*(v)\big)(a)$ for (some, and hence any) v such that $a \in \lambda(v)$. The map $\lambda^*\colon V \to Z$ is defined inductively – starting form any designated root node of T – through repeated application of the forth-property.

Lemma 4 *Let* $\pi_i\colon \mathfrak{A}_i^* \sim_G \mathfrak{A}$ *for* $i = 1, 2$ *be unravellings of* $\mathfrak{A}$*. Given guarded tuples* $\mathbf{a}_i$ *in* $\mathfrak{A}_i^*$ *such that* $\pi_1(\mathbf{a}_1) = \pi_2(\mathbf{a}_2)$*, there is an unravelling* $\pi\colon \mathfrak{A}^* \sim_G \mathfrak{A}$ *with isomorphic embeddings* $h_i\colon \mathfrak{A}_i^* \to \mathfrak{A}^*$ *such that* $h_1(\mathbf{a}_1) = h_2(\mathbf{a}_2)$ *and* $\pi_i = \pi \circ h_i$ *for* $i = 1, 2$*.*

Let T_i, λ_i be the guarded tree decompositions of $\mathfrak{A}_i^*$; u_i nodes in T_i for which $\mathbf{a}_i \subseteq \lambda_i(u_i)$ and w.l.o.g. with $\pi_1(\lambda_1(u_1)) = \pi_2(\lambda_2(u_2))$. Then we obtain a new unravelling $\mathfrak{A}^*$ by glueing $\mathfrak{A}_1^*$ and $\mathfrak{A}_2^*$ in $\mathbf{a}_1$, $\mathbf{a}_2$ as follows. Let $\mathfrak{A}^*$ be obtained from the disjoint union of $\mathfrak{A}_1^*$ and $\mathfrak{A}_2^*$ by identifying $\mathbf{a}_1$ with $\mathbf{a}_2$ via $\pi_2 \circ \pi_1^{-1}$. We let $\pi\colon \mathfrak{A}^* \to \mathfrak{A}$ be the obvious mapping induced by π_1 and π_2. Finally we obtain a tree T through glueing the T_i at the nodes u_i and letting λ be the natural merger of the λ_i over T.

3 Lifting guarded maps

Definition 5 (i) A mapping $g\colon \mathfrak{B} \to \mathfrak{A}$ between relational structures $\mathfrak{A}$ and $\mathfrak{B}$ of not necessarily the same vocabulary is *guarded* if its natural extension to subsets of B according to $s \subseteq B \mapsto \{f(b)\colon b \in s\} \subseteq A$ maps $S[\mathfrak{B}]$ into $S{\downarrow}[\mathfrak{A}]$.

(ii) If $\pi\colon \mathfrak{A}^* \sim_G \mathfrak{A}$ is a cover of $\mathfrak{A}$ and $g\colon \mathfrak{B} \to \mathfrak{A}$ a guarded mapping into $\mathfrak{A}$, then a *lift* of g to $\mathfrak{A}^*$ is a guarded mapping $g^*\colon \mathfrak{B} \to \mathfrak{A}^*$ such that $\pi \circ g^* = g$.

Homomorphisms are special cases of guarded maps.

3.1 Cycles and cliques

For the analysis of cycles we are interested in guarded mappings of n-cycles or n-cliques into $\mathfrak{A}$, $n \geqslant 3$. The n-cycle is the graph $\mathfrak{C}_n = \big(\mathbb{Z}/n\mathbb{Z}, E\big)$ with $E = \{(i, j)\colon |i - j| \equiv 1\}$. Whenever we talk about cycles we implicitly use cyclic indexing conventions and write e.g. just $(i, i+1)$ for index pairs including $(n-1, 0)$ for the appropriate n. The n-clique is the graph $\mathfrak{K}_n = \big(\{1, \ldots, n\}, E\big)$ with $E = \{(i, j)\colon i \neq j\}$.

Definition 6 Let $n \geqslant 2$.

(i) A *cycle* of length n in $\mathfrak{A}$ is a guarded mapping $g\colon \mathfrak{C}_n \to \mathfrak{A}$ such that $g(i) \neq g(i+1)$ for all $i \in \mathbb{Z}/n\mathbb{Z}$.

(ii) A *sub-cycle* of $g\colon \mathfrak{C}_n \to \mathfrak{A}$ is a cycle $h\colon \mathfrak{C}_m \to \mathfrak{A}$ for some $2 < m < n$ where $h(j) = f(i_j)$ for some strictly increasing tuple $(i_0, \ldots, i_{m-1})$ from $(0, \ldots, n-1)$.

A cycle $g\colon \mathfrak{C}_n \to \mathfrak{A}$ in $\mathfrak{A}$ is

(i) *proper* if it does not admit a lift to any guarded unravelling of $\mathfrak{A}$.

(ii) *minimal* if it does not have any sub-cycles.

Note that a cycle in $\mathfrak{A}$ is just a cycle in $\mathfrak{G}(\mathfrak{A})$ in the usual graph theoretic sense. Cycles of length 2 are degenerate and cannot be proper; however, they are (trivially) minimal. We do not consider sub-cycles of length 2 lest the notion of minimality become trivial. By the condition on $g(i) \neq g(i+1)$, cycles cannot contain loops, but they may have double points. Cycles of length 3 must involve three distinct nodes and are trivially minimal. Of course 3-cycles are the same as 3-cliques – we also speak of *triangles*.

For length greater than 3, a minimal cycle can neither have double points nor chords.

Definition 7 (i) A *clique* of size n in $\mathfrak{A}$ is an injective guarded mapping $g\colon \mathfrak{K}_n \to \mathfrak{A}$, i.e., an isomorphic embedding of $\mathfrak{K}_n$ into $\mathfrak{G}(\mathfrak{A})$.

(ii) A clique $g\colon \mathfrak{K}_n \to \mathfrak{A}$ is *unguarded* if it is not covered by a single guarded set, i.e., if $\mathrm{im}(g) \notin S{\downarrow}[\mathfrak{A}]$.

Clearly, unguarded cliques could also be characterised (in the style of proper cycles) as embeddings of $\mathfrak{K}_n$ that do not admit lifts to any guarded unravelling of $\mathfrak{A}$.

It is a simple exercise to show that a triangle is proper as a 3-cycle iff it is unguarded as a 3-clique.

Conformality of $\mathfrak{A}$ (of $\mathfrak{G}(\mathfrak{A})$) says that $\mathfrak{G}(\mathfrak{A})$ has no unguarded cliques. As mentioned above, together with chordality, this gives a well-known necessary and sufficient criterion for guarded tree decomposability. As shown below, chordality can also be replaced in this context by prohibition of (minimal) proper cycles. Recall that chordality of $\mathfrak{A}$ or $\mathfrak{G}(\mathfrak{A})$ means that every cycle of length $\ell > 3$ without double points in $\mathfrak{G}(\mathfrak{A})$ has a chord (equivalently: has a sub-cycle).

Proposition 8 *$\mathfrak{A}$ is guarded tree decomposable iff it is conformal and chordal.*

Lemma 9 *(i) Any proper cycle of $\mathfrak{A}$ is either minimal or contains a proper sub-cycle.*

(ii) A minimal cycle in $\mathfrak{A}$ is proper if, and only if, it has length greater then 3 or it is an unguarded 3-cycle (i.e., an unguarded 3-clique).

(iii) $\mathfrak{A}$ has no proper cycles if, and only if, $\mathfrak{H}(\mathfrak{A})$ is chordal and has no unguarded triangles.

(iii) immediately follows from (i) and (ii).

For (i), consider first the case that the given n-cycle $g\colon i \mapsto a_i$ has a non-degenerate chord, i.e., (a_i, a_j) an edge of $\mathfrak{G}(\mathfrak{A})$ for some $0 \leqslant i < j < n$ such that $|i-j| > 1$. We claim that at least one of the cycles $g_1 := g \restriction \{i, \ldots, j\}$ or $g_2 := g \restriction (\{0, \ldots, n-1\} \setminus \{i, \ldots, j\}$ must be proper, or else g would not be. This follows from an application of Lemma 4:

two unravellings to which these sub-cycles can be lifted could be glued in the respective representations of the guarded tuple (a_i, a_j) in such a manner that the lifts combine to yield a lift of g. Similarly, if g has double points, a decomposition into several cycles at this double point can be used.

For (ii), it is clear that a guarded triangle cannot be proper as it clearly lifts to any guarded cover.

For the converse consider a minimal cycle $g\colon i \mapsto a_i$ that is of length $\ell > 3$ or an unguarded triangle in $\mathfrak{A}$. We need to show that g is proper. Minimality, and in the case of a triangle its unguarded nature, imply that no $s \in S[\mathfrak{A}]$ can intersect $\mathrm{im}(g)$ in more than two points or in two-element sets other than the $\{a_i, a_{i+1}\}$. Assume that $\pi : \mathfrak{A}^* \sim_G \mathfrak{A}$ is an unravelling with tree decomposition $T = (V, E)$, $\lambda\colon V \to S[\mathfrak{A}^*]$. Assume $g^*\colon i \mapsto a_i^*$ is a lift of g to $\mathfrak{A}^*$. Then there are $s \in S[\mathfrak{A}^*]$ such that $a_i^*, a_{i+1}^* \in s_i$ and no $s \in S[\mathfrak{A}^*]$ intersects $\mathrm{im}(g^*)$ in other two-element sets or in larger sets. Consider $a = a_1^*$ and s_{n-1}, s_0, s_1, s_2 and let $v_{n-1}, v_0, v_1, v_2 \in V$ be such that $\lambda(v_i) = s_i$. Then v_0 is connected to v_1 in (V, E) on a path within the set $V_a = \{u\colon a \in \lambda(u)\}$. The (possibly identical) vertices v_{n-1} and v_2, on the other hand, lie outside V_a and are connected on a (possibly trivial) path outside V_a. As v_0 is also connected to v_{n-1} and v_1 to v_2 (and as at least v_0, v_1, v_2 are distinct), (V, E) would have to be cyclic, a contradiction.

If, on the other hand, $g\colon i \mapsto a_i$ is a minimal cycle that is proper and of length 3, then

As a corollary of the argument just outlined we get the following.

Corollary 10 *If $g\colon \mathfrak{C}_n \to \mathfrak{A}$ is a minimal proper cycle, then any $s \in S[\mathfrak{A}]$ either does not hit the image of g at all, or hits it in exactly one point, or hits it in two consecutive points $g(i), g(i+1)$.*

It also follows from the lemma that $\mathfrak{A}$ is guarded tree decomposable if, and only if, it has neither proper cycles nor unguarded cliques:

Assume first that $\mathfrak{A}$ has no proper cycles or unguarded cliques. Then $\mathfrak{A}$ is chordal as $\mathfrak{A}$ has no minimal cycles of length at least 4; and $\mathfrak{A}$ is conformal as $\mathfrak{A}$ has no unguarded cliques. Therefore, $\mathfrak{A}$ has a guarded tree decomposition.

Conversely, if $\mathfrak{A}$ has a guarded tree decomposition, then $\mathfrak{A}$ is its own guarded unravelling and cannot have proper cycles nor unguarded cliques.

3.2 Lifts to unravellings

We address the question, when exactly does a guarded mapping $h\colon \mathfrak{B} \to \mathfrak{A}$ lift to an unravelling of $\mathfrak{A}$? Such maps h will be regarded as unavoidable homomorphisms, because any structure that is guarded bisimilar to $\mathfrak{A}$ will admit corresponding homomorphisms.

The goal will then be the construction of covers that do not lift certain maps that fail this criterion (key case: minimal proper cycles) and eventually of covers that do not admit homomorphic images of some structure unless $\mathfrak{A}$ has homomorphic images of this structure that lift to guarded unravellings and hence to any structure that is bisimilar to $\mathfrak{A}$, too.

Definition 11 A guarded mapping $h\colon \mathfrak{B} \to \mathfrak{A}$ is *guarded tree decomposable* (or $\mathfrak{B}$ is tree decomposable in $\mathfrak{A}$ via h) if there is a chordal and conformal hypergraph $\mathfrak{H} = (B, S)$ on B such that

(i) for all $s \in S[\mathfrak{B}]$ there is some $s' \in S$ such that $s \subseteq s'$. [$S[\mathfrak{B}]$ is guarded by S.]

(ii) for all $s \in S$: $h(s) \in S{\downarrow}[\mathfrak{A}]$. [$h(S)$ is guarded by $S[\mathfrak{A}]$.]

The intuitive meaning of this criterion is that, while $\mathfrak{B}$ may not itself be guarded tree decomposable, we do achieve a tree decomposition if sets guarded by pull-backs of guarded patches of $\mathfrak{A}$ are made available as patches for the decomposition of $\mathfrak{B}$.

Proposition 12 *For any guarded mapping $h\colon \mathfrak{B} \to \mathfrak{A}$, t.f.a.e.:*

(i) h lifts to an unravelling of $\mathfrak{A}$.

(ii) h is guarded tree decomposable in $\mathfrak{A}$.

If h is guarded tree decomposable in $\mathfrak{A}$, then the tree decomposition based on subsets of h-pre-images of guarded subsets of $\mathfrak{A}$ can be used to construct a lift of h to $\mathfrak{A}^*$. Every guarded subset of $\mathfrak{A}$ can be lifted, in a manner that respects overlap with any given guarded subset in $\mathfrak{A}^*$, and the lifts of corresponding pieces of h can be glued inductively proceeding from the root to the leaves of the underlying tree.

Conversely, if h lifts to some unravelling $\mathfrak{A}^*$, then a tree decomposition of $\mathfrak{B}$ is obtained by pulling back any tree decomposition of $\mathfrak{A}^*$ in the natural manner.

Clearly, if $h\colon \mathfrak{B} \to \mathfrak{A}$ and $\pi\colon \mathfrak{A} \to \mathfrak{A}'$ are both guarded, then $\pi \circ h$ is guarded and tree decomposability of h implies tree decomposability of $\pi \circ h$.

4 Avoiding minimal proper cycles

Lemma 13 *Let $\mathfrak{A}$ be finite, $g\colon \mathfrak{C}_n \to \mathfrak{A}$ a minimal proper cycle in $\mathfrak{A}$. Then there is a finite cover $\pi\colon \mathfrak{B} \sim_G \mathfrak{A}$, which does not admit a lift of g.*

To prepare for the construction recall

$$\begin{aligned} \mathfrak{S}(\mathfrak{A}) &= \big(S[\mathfrak{A}], E\big), \\ E &= \big\{(s,s')\colon s \neq s', s \cap s' \neq \emptyset\big\}. \end{aligned}$$

Let $p\colon G = (V,E) \sim \mathfrak{S}(\mathfrak{A})$ be a bisimulation induced by the set $\{(v,p(v))\colon v \in V\}$, i.e., a *bisimilar cover*. We claim that in this case, $\mathfrak{G}$ induces a τ-structure which is a guarded cover of $\mathfrak{A}$, along the same lines that we constructed an unravelling form a tree above.

Let $\mathfrak{A}_v = \mathfrak{A} \restriction p(v)$, where $p(v) \in S[\mathfrak{A}]$ is a maximally guarded subset of $\mathfrak{A}$. Let $\mathfrak{B}'$ be the disjoint union of the $\mathfrak{A}_v$. Let the elements of $\mathfrak{B}'$ be denoted by the pairs (a,v) where $a \in \mathfrak{A}_v$, i.e. $a \in p(v)$. Consider on $\mathfrak{B}'$ the equivalence relation $\approx$ which is the reflexive transitive closure of the relation relating $a \in \mathfrak{A}_u$ to $a \in \mathfrak{A}_v$ iff $(u,v) \in E$ and $a \in \pi(u) \cap \pi(v)$. Explicitly and in general, then, we have $(a,u) \approx (a',w)$ iff $a = a'$ and there is a path in $\mathfrak{G}$ joining u to w and such that $a \in p(v)$ for all nodes v along the path.

Let $\mathfrak{B}$ be the following τ-structure on the $B'/{\approx}$, where we write $[a,v]$ for the $\approx$-class of (a,v) and $\pi\colon [a,v] \mapsto a$ for the natural projection to A. For an r-ary relation $R \in \tau$ we let $R^{\mathfrak{B}}$ be the set of all r-tuples $([a_1,v],\ldots,[a_r,v])$ in B such that $(a_1,\ldots,a_r) \in R^{\mathfrak{A}_v}$

for some $v \in V$. As $\approx$ only relates elements which correspond to one and the same element of $\mathfrak{A}$, the choice of v such that $\pi(\mathbf{a}) \in p(v)$ does not matter.

Then $\mathfrak{B}$ is a guarded cover of $\mathfrak{A}$ via $\pi \colon \mathfrak{B} \sim_G \mathfrak{A}$.

The maximally guarded subsets of $\mathfrak{B}$ are the sets $\{[a, v] \colon a \in p(v) \in S[\mathfrak{A}]\}$, and by construction the restrictions of $\pi \colon \mathfrak{B} \sim_G \mathfrak{A}$ to these patches are partial isomorphisms with $\mathfrak{A}_v$. The back-and-forth properties for this system of partial isomorphisms mirror the back-and-forth properties in the bisimulation $p \colon G = (V, E) \sim \mathfrak{S}(\mathfrak{A})$.

Consider the forth-property w.r.t. a move from the local isomorphism induced by the restriction of π to the patch $\{[a, u] \colon a \in p(u)\}$ with image $\mathfrak{A}_u$, to a patch $\{[a, v] \colon a \in p(v)\}$ in $\mathfrak{B}$. The restriction of π to the new patch has image $\mathfrak{A}_v$; this is adequate since the intersection of the u-patch and the v-patch is $\{[a, v] \colon a \in p(v) \cap p(u)\} = \{[a, u] \colon a \in p(v) \cap p(u)\}$ the pre-image of $\mathfrak{A}_u \cap \mathfrak{A}_v$, where the two restrictions of π do of course agree.

Now, for the back property w.r.t. a move from the local isomorphism induced by the restriction of π to the patch $\{[a, u] \colon a \in p(u)\}$ with image $\mathfrak{A}_u$, to some maximally guarded subset $s \in S[\mathfrak{A}]$: let v be a back-response to the move from $p(u) \in \mathfrak{S}(\mathfrak{A})$ to s, and use the restriction of π to the patch $\{[a, v] \colon a \in p(v)\}$.

We consider the fixed minimal proper cycle $g \colon \mathfrak{C}_n \to \mathfrak{A}$, $g \colon i \mapsto a_i$. Let $P = \mathbb{Z}/4n\mathbb{Z} \times S{\downarrow}[\mathfrak{C}]$, where $\mathbb{Z}/4n\mathbb{Z} = \{0, \ldots, 4n-1\}$ is regarded as a cyclic group and $S{\downarrow}[\mathfrak{C}]$ is the set of all subsets $d \subseteq \mathbb{Z}/n\mathbb{Z}$ of the form $\emptyset$, $\{i\}$, or $\{i, i+1\}$ (cyclic indexing!).

We construct $\mathfrak{G} = (V, E)$ on

$$V = \big\{(s, (m, d)) \in S[\mathfrak{A}] \times P \colon s \cap \{0, \ldots, n-1\} = d\big\}.$$

This node set is meant to record progress above $g(\mathfrak{C})$ with a 'counter' $d \in S{\downarrow}[\mathfrak{C}]$ in such a manner that the numerical counter m precludes the possibility that a cycle above $g(\mathfrak{C})$ closes in n steps.

Let $\pi \colon V \to S[\mathfrak{A}]$ be the natural projection onto the first factor, which is the universe of $\mathfrak{S}(\mathfrak{A})$. We want to interpret E on $\mathfrak{G}$ in such a way that $\pi \colon \mathfrak{G} \sim \mathfrak{S}(\mathfrak{A})$. To this end put $\big((s, (m, d)), (s', (m', d'))\big) \in E$ iff $(s, s') \in E$ in $\mathfrak{S}(\mathfrak{A})$ and one of the following conditions is satisfied

(i) $m = m'$ and $d = d'$.

(ii) $m' \equiv m + 1$ and $(d = \{i, i+1\},\ d' = \{i+1\})$ or $(d = \{i\},\ d' = \{i, i+1\})$.

(iii) $m \equiv m' + 1$ and $(d' = \{i\},\ d = \{i-1, i\})$ or $(d' = \{i, i+1\},\ d = \{i\})$.

(iv) $m' \equiv m + 2$, $d = \{i, i+1\}$, $d' = \{i+1, i+2\}$.

(v) $m \equiv m' + 2$, $d' = \{i, i+1\}$, $d = \{i+1, i+2\}$.

With this choice for E, $\mathfrak{G}$ is an undirected graph which is bisimilar to $\mathfrak{S}(\mathfrak{A})$ through $\pi \colon \mathfrak{G} \to \mathfrak{S}(\mathfrak{A})$. Let $\pi \colon \mathfrak{B} \sim_G \mathfrak{A}$ be obtained as above. We need to show that $g \colon \mathfrak{C}_n \to \mathfrak{A}$ does not lift to $\mathfrak{B}$. Suppose, towards a contradiction, that for suitable $\hat{a}_i \in \pi^{-1}(a_i)$ the mapping $\hat{g} \colon \mathfrak{C}_n \to \mathfrak{B}$, $i \mapsto \hat{a}_i$ is guarded. Choose $\hat{s}_i \in S[\mathfrak{B}]$ such that $\hat{a}_{i-1}, \hat{a}_i \in \hat{s}_i$. Let $s_i = \pi(\hat{s}_i) \in S[\mathfrak{A}]$ be the corresponding sequence in $\mathfrak{A}$; it follows that $s_i \cap \{a_0, \ldots, a_{n-1}\} = \{a_{i-1}, a_i\}$. Let further v_i be a corresponding vertex in $\mathfrak{G}$ for which $\hat{s}_i = \{[a, v_i] \colon a \in \pi(v_i)\}$. It follows that $\pi(v_i) = s_i$ and that v_i is of the form $v_i = (s_i, (m_i, d_i)) \in V$ and that $d_i = s_i \cap \{a_0, \ldots, a_{n-1}\} = \{a_{i-1}, a_i\}$.

As $\hat{a}_i \in \hat{s}_i \cap \hat{s}_{i+1}$ we know that there is an E-path $\sigma_{i+1} = u_{i0} \ldots u_{ij} \ldots u_{il}$, $u_{i0} = v_i$, $u_{il} = v_{i+1}$ in $\mathfrak{G}$ such that $a_i \in \pi(u_{ij})$ all along this path. We claim that this entails

$m_{i+1} = m_i + 2$. Let $u_{ij} = (s_{ij}, (m_{ij}, d_{ij})$, for $j = 0, \ldots, l$, so that $s_{ij} = \pi(u_{ij})$ and $d_{ij} = s_{ij} \cap \{a_0, \ldots, a_{n-1}\}$. Note that d_{ij} can only take values d_i, $\{a_i\}$, or d_{i+1}, by Corollary 10. This further implies, by induction on j, that

$$m_{ij} = \begin{cases} m_i & \text{if } d_{ij} = d_i \\ m_i + 1 & \text{if } d_{ij} = \{a_i\} \\ m_i + 2 & \text{if } d_{ij} = d_{i+1}. \end{cases}$$

Therefore, $m_{i+1} = m_i + 2$, which when considered for all i, leads to a contradiction.

We may apply the above construction of (V, E) separately for every individual minimal proper cycle of $\mathfrak{A}$. The natural product graph $(\hat{V}, \hat{E})$ (whose node set is the cartesian product of the node sets of the individual (V, E) and with edges precisely if there is an edge in every component, corresponding to a *synchronous product*) is also bisimilar to $\mathfrak{S}(\mathfrak{A})$. A structure $\mathfrak{B}$ built in the above manner on $(\hat{V}, \hat{E})$, is again a finite cover for $\mathfrak{A}$, and does not admit a lift of any minimal proper cycle of $\mathfrak{A}$.

Corollary 14 *Any finite $\mathfrak{A}$ has a finite cover $\pi\colon \mathfrak{B} \sim_G \mathfrak{A}$ that does not lift any minimal proper cycle of $\mathfrak{A}$.*

5 Avoiding homomorphic images

Consider the effect of stacking covers

$$\mathfrak{A}_2 \xrightarrow{\pi_2} \mathfrak{A}_1 \xrightarrow{\pi_1} \mathfrak{A}.$$

Observe that any unravelling of $\mathfrak{A}_2$ is an unravelling of $\mathfrak{A}_1$ and every unravelling of $\mathfrak{A}_1$ is an unravelling of $\mathfrak{A}$. A homomorphism $h\colon \mathfrak{B} \to \mathfrak{A}_2$ induces homomorphisms $\pi_2 \circ h$ into $\mathfrak{A}_1$ and $\pi_1 \circ \pi_2 \circ h$ into $\mathfrak{A}$. If one of these does lift to an unravelling (of the respective structure, and thus of those below), then $\mathfrak{B}$ is tree decomposable in $\mathfrak{A}$ via $\pi_1 \circ \pi_2 \circ h$. Tree decomposability of $\mathfrak{B}$ in $\mathfrak{A}_2$ via h implies tree decomposability in $\mathfrak{A}_1$ via $\pi_2 \circ h$ which implies tree decomposability in $\mathfrak{A}$ via $\pi_1 \circ \pi_2 \circ h$, but none of these implications can in general be reversed.

Suppose $g\colon \mathfrak{C}_n \to \mathfrak{A}_2$ is a minimal proper cycle. If the cover π_2 is as in the previous corollary, then $\pi_2 \circ g\colon \mathfrak{C}_n \to \mathfrak{A}_1$ cannot be a minimal proper cycle. In fact, it could be that $\pi_2 \circ g$ is not injective (double points violating minimality), or that the image cycle is not minimal due to the existence of chords (that are broken up in the cover), or that the image cycle is no longer proper (i.e., contains no proper sub-cycle) and hence does admit a lift to an unravelling. If the image cycle still is proper, it cannot be minimal and thus must have proper sub-cycles of length up to $n - 1$.

Stacking N levels of covers as in the previous corollary on top of a first cover $\mathfrak{A}_0$ of $\mathfrak{A}$ which is conformal (see Hodkinson/Otto [1]), we obtain a cover $\pi\colon \hat{\mathfrak{A}} = \mathfrak{A}_N \to \mathfrak{A}$. Now consider some guarded map $\hat{h}\colon \mathfrak{B} \to \hat{\mathfrak{A}}$. We claim that, if N is sufficiently large in relation to $|B|$, then the composition $h := \pi \circ \hat{h}\colon \mathfrak{B} \to \mathfrak{A}$ is guarded.

In the diagram, h_i is the guarded map obtained by composing $\hat{h}$ with the chain of the π_j up to π_i into $\mathfrak{A}_i$.

$$\begin{tikzcd}
\mathfrak{A}_N \arrow[r, "\pi_N"] & \mathfrak{A}_{N-1} \arrow[r, "\pi_{N-1}"] & \cdots \arrow[r, "\pi_{i+1}"] & \mathfrak{A}_i \arrow[r, "\pi_i"] & \cdots \arrow[r, "\pi_1"] & \mathfrak{A}_0 \arrow[r, two heads, "\pi_0"] & \mathfrak{A} \\
\mathfrak{B} \arrow[u, "\hat{h}"] \arrow[urrr, "h_i"] \arrow[urrrrrr, "h"']
\end{tikzcd}$$

Stepping down through successive levels in the stack of the $\mathfrak{A}_i$, we want to collect a chordal decomposition for all cycles in $\mathfrak{B}$. We accumulate chords which are induced by edges of $\mathfrak{G}(\mathfrak{A}_i)$ to successively augment $(B, S[\mathfrak{B}])$ towards a chordal hypergraph (B, S). Since the new edges introduced in $\mathfrak{G}(B, S)$ are h_i-pre-images of edges from some $\mathfrak{G}(\mathfrak{A}_i)$, $h_0\colon (B, S) \to \mathfrak{A}_0$ is still guarded. As $\mathfrak{A}_0$ is conformal, we may then further augment the already chordal hypergraph (B, S) to a chordal and conformal hypergraph (B, S^*) such that $\mathfrak{G}(B, S) = \mathfrak{G}(B, S^*)$ (no new edges) and such that h_0 is still guarded as a map $h_0\colon (B, S^*) \to \mathfrak{A}_0$. It follows that $h_0\colon \mathfrak{B} \to \mathfrak{A}_0$ is guarded tree decomposable. Therefore also $h\colon \mathfrak{B} \to \mathfrak{A}$ is guarded tree decomposable.

Note, however, that addition of edges induced by $\mathfrak{G}(\mathfrak{A}_i)$ to $\mathfrak{G}(B, \cdot\,)$ in the augmentation from $S[\mathfrak{B}]$ to S does create new cycles. It is therefore important that we work in B rather than on the hypergraphs induced by the $\mathfrak{A}_i$ on the h_i-images of B and that only cycles generated in a chordal decomposition of cycles of $\mathfrak{G}(\mathfrak{B})$ are successively treated until chordality is achieved. To illustrate the importance of these points, consider for instance a structure $\mathfrak{B}$ like a simple path of length 4, which is itself guarded tree decomposable. This structure $\mathfrak{B}$ may be embedded as a simple path by h_i into $\mathfrak{A}_i$ all the way down to $\mathfrak{A}_1$ before turning into a chordless cycle of length 4 in $\mathfrak{A}_0$ because π_1 identifies its end points.

To see that chordality can be achieved in the intended manner, look at a sequence of augmentations

$$S_N := S[\mathfrak{B}] \subseteq S_{N-1} \subseteq \cdots S_{i+1} \subseteq S_i \subseteq \cdots \subseteq [B]^2,$$

where $S_i \setminus S_{i+1}$ consists of edges $\{b_1, b_2\} \in [B]^2$ whose h_i-image is guarded in $\mathfrak{A}_i$ (either because $h_i(b_1) = h_i(b_2)$ or because $h_i(\{b_1, b_2\})$ is an edge of $\mathfrak{G}(\mathfrak{A}_i)$) that are non-trivial chords for cycles in $\mathfrak{G}(B, S_{i+1})$. It follows that $S_i \subseteq \big\{s \subseteq B \colon h_i(s) \in S{\downarrow}[\mathfrak{A}]\big\}$ so that $h_i\colon (B, S_i) \to \mathfrak{A}_i$ remains guarded. As any cycle in $\mathfrak{G}(B, S_{i+1})$ whose π_i-image in $\mathfrak{G}(\mathfrak{A}_i)$ would still contain a chordless sub-cycle of length greater than 3 would be a lift of a minimal proper cycle of $\mathfrak{A}_i$ to its guarded cover $\pi_i\colon \mathfrak{A}_{i+1} \to \mathfrak{A}_i$, we know that any chordless cycle gets eliminated within a number of steps that is bounded by $|B|$ (the maximal length cycles to be considered). In this manner, $S_{i+1} = S_i$ only when (B, S_{i+1}) is already chordal, in which case we let $S := S_{i+1}$. By monotonicity of the sequence of the S_i, ${}_i = Si + 1$ for some $i < N - |[B]^2|$. It follows that we find a set S as desired whenever $N \geqslant |B|^2$.

Hence $\hat{h}\colon \mathfrak{B} \to \mathfrak{A}_N$ is guaranteed to be guarded tree decomposable in $\mathfrak{A}_N$ if $|B|^2 \leqslant N$.

Theorem 15 *Let $\mathfrak{A}$ be finite, $n \geqslant 3$. Then there is a finite cover $\pi\colon \hat{\mathfrak{A}} \sim_G \mathfrak{A}$ that has homomorphic images of structures of size up to n only if these are unavoidable in the* *sense that every unravelling of $\mathfrak{A}$ admits such homomorphic images. More precisely: if $\hat{h}\colon \mathfrak{B} \to \hat{\mathfrak{A}}$ is a homomorphism with $|\mathfrak{B}| \leqslant n$, then the combined homomorphism $h := \pi \circ \hat{h}$ lifts to any unravelling of $\mathfrak{A}$.*

Corollary 16 *Let φ be a sentence of* $\mathrm{GF}[\tau]$*, $\mathfrak{B}$ a finite τ-structure. If φ has a model which does not homomorphically embed $\mathfrak{B}$ then φ has a finite model not embedding $\mathfrak{B}$.*

Suppose every finite model of φ admits a homomorphic image of $\mathfrak{B}$. Then, by the theorem, every finite model of $\mathfrak{B}$ admits a homomorphic image of $\mathfrak{B}$ that lifts to all unravellings. It follows that every finite model has a guarded tree decomposable homomorphic image of $\mathfrak{B}$. So for finite, and hence all structures, we have the valid implication

$$\varphi \models \psi_{\mathfrak{B}},$$

where $\psi_{\mathfrak{B}}$ is the finite disjunction of the existential guarded formulae describing every manner in which homomorphic images of $\mathfrak{B}$ can be tree decomposable in any structure $\mathfrak{A}$. But this means that φ cannot have an infinite model that avoids homomorphic images of $\mathfrak{B}$ either.

References

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