

Durham Research Online

Deposited in DRO:

06 April 2010

Version of attached file:

Published Version

Peer-review status of attached file:

Peer-reviewed

Citation for published item:

Dalmau, V. and Krokhin, A. and Larose, B. (2004) 'First-order definable retraction problems for posets and reflexive graphs.', in 19th annual IEEE symposium on logic in computer science, LICS'04, 13-17 July 2004, Turku, Finland ; proceedings. , pp. 232-241.

Further information on publisher's website:

<http://dx.doi.org/10.1109/LICS.2004.1319617>

Publisher's copyright statement:

©2004 IEEE. This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Additional information:

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in DRO
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full DRO policy](#) for further details.

First-order definable retraction problems for posets and reflexive graphs*

Víctor Dalmau
Department of Technology
University Pompeu Fabra
Barcelona 08003, Spain
victor.dalmau@upf.edu

Andrei Krokhin
Department of Computer Science
University of Durham
Durham, DH1 3LE, UK
andrei.krokhin@durham.ac.uk

Benoit Larose
Department of Mathematics and Statistics
Concordia University
Montréal, Canada H3G 1M8
larose@mathstat.concordia.ca

November 10, 2005

Abstract

A retraction from a structure \mathbf{P} to its substructure \mathbf{Q} is a homomorphism from \mathbf{P} onto \mathbf{Q} that is the identity on \mathbf{Q} . We present an algebraic condition which completely characterises all posets and all reflexive graphs \mathbf{Q} with the following property: the class of all posets or reflexive graphs, respectively, that admit a retraction onto \mathbf{Q} is first-order definable.

Keywords: retraction, homomorphism, graphs, posets, first-order definability

1 Introduction

Let π be a vocabulary that may contain only relation and constant symbols. Throughout the paper we use the same boldface and slanted capital letters to denote a structure and its universe, respectively. Recall that a *homomorphism* from a π -structure \mathbf{P} to a π -structure \mathbf{Q} is a mapping h from P to Q such that, for any relation symbol R in π , we have $h(\bar{x}) \in R^{\mathbf{Q}}$ whenever $\bar{x} \in R^{\mathbf{P}}$, and, for any constant symbol $c \in \pi$, we have $h(c^{\mathbf{P}}) = c^{\mathbf{Q}}$. If, in addition, \mathbf{Q} is a substructure of \mathbf{P} and h fixes every element of Q then h is said to be a *retraction*.

Homomorphism and retraction problems have been an object of intensive study in combinatorics, logic, and computer science. The homomorphism problem for a fixed π -structure \mathbf{Q} (denoted $Hom(\mathbf{Q})$) is whether a given π -structure \mathbf{P} admits a homomorphism to \mathbf{Q} . The retraction problem for \mathbf{Q} (denoted $Ret(\mathbf{Q})$) is defined similarly. The homomorphism and retraction problems are equivalent to *constraint satisfaction problems* that are much studied in computer science (and, in particular, in finite model theory) and artificial intelligence (see, e.g., [5, 7, 14, 33]).

Note that $Hom(\mathbf{Q})$ and $Ret(\mathbf{Q})$ can be viewed as classes of structures that admit a homomorphism or retraction, respectively, to \mathbf{Q} . Hence one can try to describe these classes (or their complements $\neg Hom(\mathbf{Q})$

* A short version of this paper appeared in the Proceedings of the 19th Symposium on Logic in Computer Science (LICS 2004).

and $\neg Ret(\mathbf{Q})$ in various logics (see, e.g., [1, 8, 9, 14, 22]). In this paper, we will describe the complements because they are homomorphism-closed and it is perhaps more customary to consider homomorphism-closed classes (see, e.g., [14, 15]). If \mathbf{Q} is a structure over the relational vocabulary π , and $Q = \{q_1, \dots, q_k\}$ then, in order to describe $Ret(\mathbf{Q})$ or $\neg Ret(\mathbf{Q})$ in logics, it is natural to consider structures over the vocabulary $\sigma = \pi \cup \{c_1, \dots, c_k\}$ obtained from π by adding k constant symbols c_1, \dots, c_k , viewing \mathbf{Q} as a σ -structure such that $c_i^{\mathbf{Q}} = q_i$ for all i , and assuming that *all* σ -structures \mathbf{P} under consideration *contain* \mathbf{Q} , that is, the constants interpret in \mathbf{P} as a substructure isomorphic to \mathbf{Q} under the map $c_i^{\mathbf{P}} \mapsto q_i$. A closely related problem to $Ret(\mathbf{Q})$ is the *one-or-all list homomorphism* problem which can be viewed as the homomorphism problem, where, in addition, each element of the input structure is assigned a list of possible target values in \mathbf{Q} , and each of these lists consists either of a single element or of Q itself. More formally, let $\tau = \pi \cup \{R_1, \dots, R_k\}$ where R_1, \dots, R_k are unary relation symbols, and view \mathbf{Q} as a τ -structure with $R_i^{\mathbf{Q}} = \{q_i\}$ for all i . We denote by $Hom_{\tau}(\mathbf{Q})$ the class of all τ -structures \mathbf{P} admitting a homomorphism to \mathbf{Q} (as a τ -structure). Such problems have been considered for graphs (see, e.g., [13]). It is easy to notice that from the computational complexity point of view $Ret(\mathbf{Q})$ and $Hom_{\tau}(\mathbf{Q})$ are always equivalent. However, it is not so evident, if true at all, that these problems have the same descriptive complexity (that is, for any logic L , they are definable or not definable in L simultaneously).

In this paper we will consider two important special cases of problems $Ret(\mathbf{Q})$ and $Hom_{\tau}(\mathbf{Q})$: those where π consists of a single binary relation symbol and \mathbf{Q} is either a poset or a reflexive (undirected) graph. (Note that $Hom(\mathbf{Q})$ for such structures is trivial). Retractions play an important role in the structure theory of graphs and orders [11, 17, 18]. The computational complexity of these problems has been extensively studied [13, 14, 16, 23, 27, 31] in an attempt to distinguish tractable cases from NP-complete ones. However, this seems to be a very difficult problem in general because it is known [13, 14] that every constraint satisfaction problem can be encoded as $Ret(\mathbf{Q})$ for a suitable poset or for a suitable reflexive graph \mathbf{Q} , and general constraint satisfaction problems are known to be very difficult to classify. Here we will study the descriptive complexity of these problems: classify problems $Ret(\mathbf{Q})$ and $Hom_{\tau}(\mathbf{Q})$ (for all possible \mathbf{Q}) with respect to definability in a given logic L . More specifically, we take L to be the most studied logic **FO** (first-order logic), and we give a *complete* classification for this case. Atserias [1] characterized all first-order definable problems of the form $Hom(\mathbf{Q})$, where \mathbf{Q} is any finite structure, as those having finitary duality (see also [29]). The results of [1] concern a larger class of problems than the one considered in this paper, but our characterization is different and more explicit – in particular, it implies that reflexive graphs and posets \mathbf{Q} for which problems $Ret(\mathbf{Q})$ and $Hom_{\tau}(\mathbf{Q})$ are first-order definable can be recognized in polynomial time. We remark that our results are also similar in spirit to the classification of Fixed Subgraph Homeomorphism problems for directed graphs with respect to definability in logical program language Datalog(\neq) and in the infinitary logic $L_{\infty, \omega}^{\omega}(\neq)$ [20, 21].

Our proofs are based on algebraic and combinatorial characterisations of certain graphs and posets [24, 25, 26] and on the Ehrenfeucht-Fraïssé method for proving inexpressibility in **FO** [12, 20]. An n -ary operation f on a relational structure \mathbf{Q} is said to be a *polymorphism* of \mathbf{Q} if it is a homomorphism from the Cartesian power \mathbf{Q}^n to \mathbf{Q} . For posets and graphs, this means that if $\bar{a}, \bar{b} \in Q^n$ are such that $(a_i, b_i) \in \theta$ for all i (where θ is the relation in \mathbf{Q}) then we also have $(f(\bar{a}), f(\bar{b})) \in \theta$. If f is a polymorphism of a poset then f is also said to be *monotone* on it. We now define operations that play a most important role in this paper. An n -ary ($n \geq 3$) operation f satisfying the condition that, for any a , $f(x_1, \dots, x_n) = a$ whenever at least $n - 1$ of the x_i 's are equal to a is called a *near-unanimity* (NU) operation. It is known that, for relational structures with a NU polymorphism, the problem $Hom(\mathbf{Q})$ can be solved in polynomial time [14, 19]. For the sake of brevity, we shall call posets or graphs with a NU polymorphism *NU-posets* and *NU-graphs*. Along with connectedness, this will be the property responsible for **FO**-definability (see Theorems 3.1 and 4.1). It was shown in [24, 25, 26] that for connected posets and reflexive graphs, this property is equivalent on the one hand to the finiteness of obstructions for $Hom_{\tau}(\mathbf{Q})$, which we use to prove **FO**-definability, and

on the other hand to connectedness of certain substructures in powers of \mathbf{Q} , which we use together with the Ehrenfeucht-Fraïssé method to show that the NU property is also a necessary condition for **FO**-definability.

The NU property has attracted much attention in algebra and combinatorics (see, e.g., [4, 10, 19, 24, 25, 26, 34]). Examples of reflexive NU-graphs include all chordal reflexive graphs while no reflexive cycle of length at least 4 has a NU polymorphism [4]. A poset is called a lattice if every pair of its elements has a least upper bound and a greatest lower bound. Lattices are the simplest examples of NU-posets. Examples of posets without NU polymorphisms are non-dismantlable posets, including all posets in which every non-minimal element covers at least two elements, and every non-maximal element is covered by at least two elements (e.g., crowns in Fig. 2) [26]. A combinatorial characterisation of (reflexive) NU graphs and NU posets was obtained in [25, 26] (see Theorem 2.1).

Note that it is not possible to bound the arity of NU polymorphisms to define NU-graphs or NU-posets, since by results of [4, 10], for every $n \geq 3$, there is a NU-poset and a reflexive NU-graph having no NU polymorphisms of arity less than n . On the positive note, NU-graphs and NU-posets can be recognized in polynomial time (in the size of the structure) [24, 25].

2 Preliminaries

Fix a structure $\mathbf{Q} = \langle Q; \theta \rangle$ where $Q = \{q_1, \dots, q_k\}$ and θ is a binary relation on Q . There are two vocabularies that we consider throughout:

- $\sigma = \{E, c_1, \dots, c_k\}$ consisting of one binary relation symbol E and k constant symbols c_1, \dots, c_k ; and
- $\tau = \{E, R_1, \dots, R_k\}$ consisting of one binary relation symbol E and k unary relation symbols R_1, \dots, R_k .

The structure \mathbf{Q} will also be interpreted both as a σ -structure and a τ -structure, where $E^{\mathbf{Q}}$ is θ , $c_i^{\mathbf{Q}}$ is q_i and $R_i^{\mathbf{Q}}$ is $\{q_i\}$ for all $1 \leq i \leq k$. It will always be clear from the context which vocabulary is assumed. Similarly, any σ -structure \mathbf{P} can be viewed as a τ -structure with $R_i^{\mathbf{P}} = \{c_i^{\mathbf{P}}\}$.

As mentioned above, we consider the following two cases for \mathbf{Q} : (A) \mathbf{Q} is a *graph*, which we assume to be reflexive (i.e. with all loops) and symmetric (i.e. undirected) or (B) \mathbf{Q} is a *poset*, i.e. where θ is reflexive, anti-symmetric and transitive.

A σ -structure \mathbf{P} is said to *contain* \mathbf{Q} if the constants interpret in \mathbf{P} as a substructure isomorphic to \mathbf{Q} under the map $c_i^{\mathbf{P}} \mapsto q_i$.

Our results will rely on the following concept of *obstruction* (also called *zig-zag* in the case of posets) introduced in [34]. We define a partial ordering on τ -structures as follows: we write that $\mathbf{H} \preceq \mathbf{H}'$ if (i) $H \subseteq H'$, (ii) $E^{\mathbf{H}} \subseteq E^{\mathbf{H}'}$ and (iii) $R_i^{\mathbf{H}} = R_i^{\mathbf{H}'} \cap H$ for all $1 \leq i \leq k$. Clearly if $\mathbf{H} \preceq \mathbf{H}'$ then the inclusion map is a homomorphism from \mathbf{H} to \mathbf{H}' .

An *obstruction* for the graph (resp. poset) \mathbf{Q} is a τ -structure \mathbf{H} where \mathbf{H} is a graph (resp. poset) such that (1) there is no homomorphism (of τ -structures) from \mathbf{H} to \mathbf{Q} , (2) \mathbf{H} is minimal with respect to property (1) (in the ordering \preceq) and (3) the unary relations in \mathbf{H} are pairwise disjoint. It is clear that for any τ -structure \mathbf{P} which is a graph or poset, there is no homomorphism from \mathbf{P} to \mathbf{Q} if and only if some unary relations of \mathbf{P} intersect or $\mathbf{H} \preceq \mathbf{P}$ for some obstruction \mathbf{H} for \mathbf{Q} . Note that there are other notions of obstruction used in the study of homomorphisms (see, e.g., [8]).

We shall also need the following notion from universal algebra: let \mathbf{Q} be a graph (resp. poset) and let $n \geq 1$: \mathbf{Q}^n denotes the Cartesian power of \mathbf{Q} , that is, $(\bar{a}, \bar{b}) \in E^{\mathbf{Q}^n}$ if and only if $(a_i, b_i) \in \theta$ for all $1 \leq i \leq n$. An *idempotent subalgebra* of \mathbf{Q}^n is a subset $X \subseteq \mathbf{Q}^n$ that can be described as follows: there exists a triple $(\mathbf{Y}, (y_1, \dots, y_n), \gamma)$ where \mathbf{Y} is a graph (resp. poset), $y_i \in Y$ for all $i = 1, \dots, n$ and

γ is a partial map from Y to Q with domain Y' with the following property: X consists of all n -tuples $(\delta(y_1), \dots, \delta(y_n))$ where $\delta : Y \rightarrow Q$ runs through all edge-preserving (resp. monotone) maps whose restriction to Y' is equal to γ .

The name ‘‘idempotent subalgebra’’ comes from an equivalent algebraic description which goes as follows. An m -ary operation f on \mathbf{Q} is called *idempotent* if $f(x, \dots, x) = x$ for all $x \in Q$. A subset $X \subseteq Q^n$ is an idempotent subalgebra of \mathbf{Q}^n if and only if it is preserved by all idempotent polymorphisms of \mathbf{Q} , that is, $f(\bar{a}_1, \dots, \bar{a}_m) \in X$ for all m -ary idempotent polymorphisms f of \mathbf{Q} and all $\bar{a}_1, \dots, \bar{a}_m \in X$, $m \geq 1$, where f acts on tuples componentwise.

Another equivalent definition for idempotent subalgebras is that they are exactly those subsets X of Q^n , $n \geq 1$, which can be defined by primitive positive first-order formulas (with n free variables) in Q (as a τ -structure). Equivalence of the last two definitions follows from Theorems 1.2.3 and 2.1.3 in [30], while equivalence of the first and the third definitions is rather straightforward.

A combinatorial characterisation of NU posets and NU graphs is based on the notion of dismantling. A graph \mathbf{G}_1 is said to *dismantle* to its subgraph \mathbf{G}_2 if there is a sequence $\mathbf{G}'_1, \dots, \mathbf{G}'_n$ of subgraphs of \mathbf{G}_1 such that $\mathbf{G}'_1 = \mathbf{G}_1$, $\mathbf{G}'_n = \mathbf{G}_2$, and, for $i = 1, \dots, n-1$, \mathbf{G}'_{i+1} is obtained from \mathbf{G}'_i by removing a vertex v_i such that v_i is dominated by some other vertex u_i in \mathbf{G}'_i (i.e., every neighbour of v_i , including v_i itself, in \mathbf{G}_i is also a neighbour of u_i). Similarly, a poset \mathbf{P}_1 is said to dismantle to its subposet \mathbf{P}_2 if there is a sequence $\mathbf{P}'_1, \dots, \mathbf{P}'_n$ of subposets of \mathbf{P}_1 such that $\mathbf{P}'_1 = \mathbf{P}_1$, $\mathbf{P}'_n = \mathbf{P}_2$, and, for $i = 1, \dots, n-1$, \mathbf{P}'_{i+1} is obtained from \mathbf{P}'_i by removing a vertex v_i such that v_i is dominated by some other vertex u_i in the comparability graph of \mathbf{P}'_i . If \mathbf{Q} is a graph or a poset, then the *diagonal* of \mathbf{Q}^2 is its subgraph (subposet, respectively) induced by all nodes of the form (x, x) . It is easy to see that the diagonal of \mathbf{Q}^2 is isomorphic to \mathbf{Q} .

The following result was proved in [24, 26] for posets and in [25] for graphs.

Theorem 2.1 *Let \mathbf{Q} be a connected graph or poset. Then the following are equivalent:*

1. *there are finitely many obstructions for \mathbf{Q} ;*
2. *\mathbf{Q} admits an NU polymorphism (of some arity);*
3. *for every $n \geq 1$, every idempotent subalgebra of \mathbf{Q}^n is connected;*
4. *\mathbf{Q}^2 dismantles to its diagonal.*

Note that, for posets, it is enough to take to $n = 1$ in condition 3 of the above theorem [26], but, for graphs, we need to consider all $n \geq 1$ [25].

We also need the following: if \mathbf{P} is a σ - or τ -structure, let $\bar{\mathbf{P}}$ denote its symmetric, reflexive closure, i.e. the similar structure on the same universe such that $(x, x) \in E^{\bar{\mathbf{P}}}$ for all $x \in P$ and $(x, y) \in E^{\bar{\mathbf{P}}}$ whenever $(x, y) \in E^{\mathbf{P}}$ or $(y, x) \in E^{\mathbf{P}}$. The following is immediate: there is a homomorphism from \mathbf{P} to the graph \mathbf{Q} if and only if there is a homomorphism from $\bar{\mathbf{P}}$ to \mathbf{Q} ; and if there is a homomorphism from \mathbf{P} to \mathbf{P}' then there is one from $\bar{\mathbf{P}}$ to $\bar{\mathbf{P}}'$.

Finally, we define for each τ -structure a sentence in the language of τ which will encode the given structure, see [6] and [9]: let \mathbf{H} be a τ -structure and let $H = \{h_1, \dots, h_l\}$. Denote by $T_{\mathbf{H}}$ the τ -sentence

$$\exists x_{h_1} \dots \exists x_{h_l} (\bigwedge E(x_h, x_{h'}) \wedge (\bigwedge R_i(x_h))$$

where the first conjunction is taken over all pairs $(h, h') \in E^{\mathbf{H}}$ and the second over all $h \in R_i^{\mathbf{H}}$ and all $1 \leq i \leq k$.

When the structure has the property that its unary relations are pairwise disjoint and each contain at most one element, we can define a sentence in the language of σ which will encode the given structure: let \mathbf{H} be such a τ -structure. Denote by H' the set of all $h \in H$ such that $h \notin R_i^{\mathbf{H}}$ for all i . For every $h \in H$, set

$h_f = x_h$ if $h \in H'$ and $h_f = c_i$, where $h \in R_i^{\mathbf{H}}$ otherwise. Suppose $H' = \{h_1, \dots, h_s\}$. Denote by $S_{\mathbf{H}}$ the σ -sentence

$$\exists x_{h_1} \dots \exists x_{h_s} (\bigwedge E(h_f, h'_f))$$

where the conjunction is taken over all pairs (h, h') such that $(h, h') \in E^{\mathbf{H}}$.

The first statement in the next lemma is a special case of a result in [6]. Both proofs are straightforward.

Lemma 2.2 *Let \mathbf{H} and \mathbf{P} be τ -structures and let \mathbf{P}' be a σ -structure.*

1. *the sentence $T_{\mathbf{H}}$ is true in \mathbf{P} if and only if there is a homomorphism from \mathbf{H} to \mathbf{P} ;*
2. *if $c_i^{\mathbf{P}'} \neq c_j^{\mathbf{P}'}$ whenever $q_i \neq q_j$, and $R_i^{\mathbf{H}}$ contains at most one element for all $1 \leq i \leq k$, then the sentence $S_{\mathbf{H}}$ is true in \mathbf{P}' if and only if there is a τ -homomorphism from \mathbf{H} to \mathbf{P}' .*

3 Graph retraction problems in FO

In this section \mathbf{Q} is a graph. Recall that by a graph we always mean a reflexive graph. Let $\neg Ret(\mathbf{Q})$ denote the class of all σ -structures that contain \mathbf{Q} but do not retract onto \mathbf{Q} . Let $\neg Hom_{\tau}(\mathbf{Q})$ denote the class of all τ -structures that do not admit a τ -homomorphism to \mathbf{Q} .

Our main result for graphs is the following:

Theorem 3.1 *Let \mathbf{Q} be a graph. Then the following conditions are equivalent:*

- (1) *the class $\neg Ret(\mathbf{Q})$ is FO-definable;*
- (2) *the class $\neg Hom_{\tau}(\mathbf{Q})$ is FO-definable;*
- (3) *\mathbf{Q} is a connected NU-graph.*

Moreover, if any of these conditions holds then both of the above classes can be defined by a first-order formula that contains neither negation nor universal quantification.

Note that the last statement of the theorem can also be obtained by combining the first part of the theorem with the Finite Homomorphism Preservation Theorem [32] recently proved by Rossman.

We shall require the following notion from finite model theory (see [12] Definition 2.3.1):

Definition. Let \mathbf{A} and \mathbf{A}' be two structures over the same vocabulary, and let m be a non-negative integer. We say that \mathbf{A} and \mathbf{A}' are said to be m -isomorphic if there exists a sequence $(I_j)_{j \leq m}$ with the following properties:

1. Every I_j is a non-empty set of partial isomorphisms from \mathbf{A} to \mathbf{A}' ;
2. (*Forth property*) For every $j < m$, $p \in I_{j+1}$ and $a \in A$ there is a $q \in I_j$ such that q is an extension of p and a is in the domain of q ;
3. (*Back property*) For every $j < m$, $p \in I_{j+1}$ and $b \in A'$ there is a $q \in I_j$ such that q is an extension of p and b is in the range of q .

The following result is well known (see Theorem 2.2.12 and Corollary 2.3.4 of [12]).

Proposition 3.2 *Let \mathcal{K} be a class of finite structures such that, for every $m \geq 1$, there are m -isomorphic structures \mathbf{A}_m and \mathbf{A}'_m with $\mathbf{A}_m \in \mathcal{K}$ and $\mathbf{A}'_m \notin \mathcal{K}$. Then \mathcal{K} is not FO-definable.*

The main technical result used in the proof is the following lemma:

Proposition 3.3 *Let \mathbf{Q} be a graph. If \mathbf{Q} is not connected or admits no NU polymorphism then, for every integer $m \geq 1$, there exist graphs \mathbf{P} and \mathbf{P}' with the following properties:*

- (i) \mathbf{P} and \mathbf{P}' contain \mathbf{Q} ;
- (ii) \mathbf{P} and \mathbf{P}' are m -isomorphic (as σ - or τ -structures);
- (iii) \mathbf{P} retracts onto \mathbf{Q} but \mathbf{P}' does not.

The proof of Proposition 3.3 is found in Section 5.

Proof of Theorem 3.1. It follows from Propositions 3.2 and 3.3 that if \mathbf{Q} is either disconnected or does not admit an NU polymorphism then none of the two classes are **FO**-definable. So now assume that \mathbf{Q} is connected and admits an NU polymorphism. First we show that (2) holds. We build a sentence as follows: by Theorem 2.1, there exist finitely many obstructions for \mathbf{Q} , say $\mathbf{H}_1, \dots, \mathbf{H}_m$. Note that the \mathbf{H}_i 's are obstructions in the class of all reflexive graphs, and general τ -structures are not necessarily reflexive graphs (with added unary relations), hence we consider the following sets. For every $1 \leq i \leq m$, let $\mathcal{L}(\mathbf{H}_i)$ denote the set of all τ -structures \mathbf{G} such that $\overline{\mathbf{G}} = \mathbf{H}_i$. Consider the sentence

$$\left(\bigvee_{i=1}^m \bigvee_{\mathbf{G} \in \mathcal{L}(\mathbf{H}_i)} T_{\mathbf{G}} \right) \vee \left(\exists x \left(\bigvee_{1 \leq i \neq j \leq k} (R_i(x) \wedge R_j(x)) \right) \right),$$

where $T_{\mathbf{G}}$ is as defined in Section 2. We claim that this sentence captures precisely those τ -structures that admit no homomorphism to \mathbf{Q} . Indeed, suppose that the sentence is true in a structure \mathbf{P} . If the second part of this sentence is true in \mathbf{P} then trivially there is no homomorphism from \mathbf{P} to \mathbf{Q} . Assume that the second part is false. This means that $T_{\mathbf{G}}$ is true in \mathbf{P} for some structure \mathbf{G} such that $\overline{\mathbf{G}} = \mathbf{H}_i$ for some i . Hence by Lemma 2.2 there exists a τ -homomorphism from \mathbf{G} to \mathbf{P} , and consequently a homomorphism from \mathbf{H}_i to $\overline{\mathbf{P}}$. It follows that we cannot have a τ -homomorphism from \mathbf{P} to \mathbf{Q} . Conversely, suppose that there is no homomorphism from \mathbf{P} to \mathbf{Q} and that the unary relations in \mathbf{P} are pairwise disjoint; thus there is no homomorphism from $\overline{\mathbf{P}}$ to \mathbf{Q} and hence there exists some i such that $\mathbf{H}_i \preceq \overline{\mathbf{P}}$. We show that there exists some $\mathbf{G} \in \mathcal{L}(\mathbf{H}_i)$ that admits a homomorphism in \mathbf{P} . Indeed, we create a substructure \mathbf{G} of \mathbf{H}_i by setting $G = H_i$, $E^{\mathbf{G}} = E^{\mathbf{H}_i} \cap E^{\mathbf{P}}$ and $R_j^{\mathbf{G}} = R_j^{\mathbf{H}_i}$ for all $1 \leq j \leq k$. Clearly \mathbf{G} is a substructure of \mathbf{P} and furthermore $\overline{\mathbf{G}} = \mathbf{H}_i$, since for each edge $(x, y) \in E^{\mathbf{H}_i}$ with $x \neq y$, we have that either (x, y) or (y, x) is in $E^{\mathbf{P}}$. Thus the sentence $T_{\mathbf{G}}$ is true in \mathbf{P} .

Next we show that (1) holds. Let $\mathcal{L}'(\mathbf{H}_i)$ denote the set of all τ -structures \mathbf{G} such that $\overline{\mathbf{G}} = \mathbf{H}_i$, and such that each $R_i^{\mathbf{G}}$ contains at most one element of G . Consider the sentence

$$\bigvee_{i=1}^m \bigvee_{\mathbf{G} \in \mathcal{L}'(\mathbf{H}_i)} S_{\mathbf{G}},$$

where $S_{\mathbf{G}}$ is as defined in Section 2. We claim that a σ -structure \mathbf{P} containing \mathbf{Q} does not retract onto \mathbf{Q} if and only if this sentence is true in \mathbf{P} . Suppose that $S_{\mathbf{G}}$ is true in \mathbf{P} for some τ -structure $\mathbf{G} \in \mathcal{L}'(\mathbf{H}_i)$. By Lemma 2.2 there is a τ -homomorphism from \mathbf{G} to \mathbf{P} , and hence from \mathbf{H}_i to $\overline{\mathbf{P}}$; consequently there cannot be a τ -homomorphism from $\overline{\mathbf{P}}$ to \mathbf{Q} ; and hence no retraction either.

For the converse, suppose that there is no retraction of \mathbf{P} onto \mathbf{Q} , i.e. there is no σ -homomorphism from \mathbf{P} to \mathbf{Q} ; then there is no τ -homomorphism either, and so there is no τ -homomorphism from $\overline{\mathbf{P}}$ to \mathbf{Q} . Now proceed exactly as in the proof of (2): there exists an obstruction \mathbf{H}_i of \mathbf{Q} such that $\mathbf{H}_i \preceq \overline{\mathbf{P}}$; and simply notice that, since \mathbf{P} is a σ -structure, the τ -structure \mathbf{G} produced as above will have at most one element in each unary relation. ■

4 Poset retraction problems in FO

In this section \mathbf{Q} is a poset. Let $\neg PoRet(\mathbf{Q})$ denote the class of all posets that contain \mathbf{Q} that do not retract onto \mathbf{Q} . Let $\neg PoHom_\tau(\mathbf{Q})$ denote the class of all posets that do not admit a τ -homomorphism to \mathbf{Q} .

We consider $PoRet(\mathbf{Q})$ rather than $Ret(\mathbf{Q})$ for the following reason: $Ret(\mathbf{Q})$ is *not* FO-definable for any poset \mathbf{Q} with more than one element. Indeed, it follows from the proof of Theorem 4.1 that if \mathbf{Q} is not connected then $Ret(\mathbf{Q})$ is not FO-definable. Assume that \mathbf{Q} is connected and fix elements $a < b$ in \mathbf{Q} such that a is minimal and b is maximal. One can slightly modify the construction in Exercise 2.3.9 of [12] to obtain m -isomorphic structures with the desired properties. Let \mathbf{P}_n be the structure obtained from \mathbf{Q} by adding to $E^{\mathbf{Q}}$ pairs forming two sufficiently long oriented (reflexive) paths $a_n, \dots, a_1 = a$ and $b = b_1, \dots, b_n$ and also add a sufficiently long oriented (reflexive) cycle C not connected with the rest of the digraph. Let \mathbf{P}'_n be the structure obtained from \mathbf{P}_n by adding the pair (b_n, a_n) to $E^{\mathbf{P}_n}$ and removing one arc from the cycle C . Clearly, \mathbf{P}_n retracts onto \mathbf{Q} (by sending all a_i 's to a and all b_i 's and C to b) while \mathbf{P}'_n does not (because it contains a cycle with a and b in it). One can show that, for every $m \geq 0$, there is a sufficiently large n such that \mathbf{P}_n and \mathbf{P}'_n are m -isomorphic, which implies that, by Proposition 3.2, $Ret(\mathbf{Q})$ is not FO-definable.

The problem $PoRet(\mathbf{Q})$ has been studied in connection with type reconstruction [2, 28] and constraint satisfaction [14, 23], and is a natural choice for a restriction of $Ret(\mathbf{Q})$.

Our main result for posets is the following:

Theorem 4.1 *Let \mathbf{Q} be a poset. The following conditions are equivalent:*

1. *the class $\neg PoRet(\mathbf{Q})$ is FO-definable;*
2. *the class $\neg PoHom_\tau(\mathbf{Q})$ is FO-definable;*
3. *\mathbf{Q} is a connected NU-poset.*

Moreover, if any of these conditions holds then both of the above classes can be defined by a first-order formula that contains neither negation nor universal quantification.

Similarly to Theorem 3.1, the last statement of the theorem can also be obtained by combining the first part of the theorem with the Finite Homomorphism Preservation Theorem.

The proof of the result is similar to the graph case, with two notable differences: first, the glueing construction we use in the poset case is more involved than in the case of graphs, because we must guarantee that the resulting structure is transitive. On the other hand, it has been shown in [24] that if a connected poset \mathbf{Q} admits no NU polymorphism then it has a disconnected idempotent subalgebra, so we do not need to consider higher powers of \mathbf{Q} (this result is not valid for graphs, see [25]). The rest of the proof of Theorem 4.1 is quite similar to that of Theorem 3.1 (and is in fact slightly simpler because there is no need to consider symmetric reflexive closures of relations). The main technical result used in the proof is the following proposition, proven in Section 6.

Proposition 4.2 *Let \mathbf{Q} be a poset. If \mathbf{Q} is not connected or admits no NU polymorphism then, for every integer $m \geq 1$, there exist posets \mathbf{P} and \mathbf{P}' with the following properties:*

1. *\mathbf{P} and \mathbf{P}' contain \mathbf{Q} ;*
2. *\mathbf{P} and \mathbf{P}' are m -isomorphic (as σ - or τ -structures);*
3. *\mathbf{P} retracts onto \mathbf{Q} but \mathbf{P}' does not.*

Proof of Theorem 4.1. Just as in Theorem 3.1 it is immediate from Lemma 4.2 that if \mathbf{Q} is either disconnected or admits no NU polymorphism then $\neg PoRet(\mathbf{Q})$ and $\neg PoHom_\tau(\mathbf{Q})$ are not FO-definable. So now assume that \mathbf{Q} is connected and admits an NU polymorphism. First we show that (2) holds. We build a sentence as follows: by Theorem 2.1, there exist finitely many obstructions for \mathbf{Q} , say $\mathbf{H}_1, \dots, \mathbf{H}_m$. Consider the sentence

$$\left(\bigvee_{i=1}^m T_{\mathbf{H}_i}\right) \vee \left(\exists x \left(\bigvee_{1 \leq i \neq j \leq k} (R_i(x) \wedge R_j(x))\right)\right).$$

We claim that this sentence captures precisely those posets that admit no τ -homomorphism to \mathbf{Q} . Indeed, suppose that the sentence is true in the structure \mathbf{P} . If the second part of this sentence is true then trivially there is no homomorphism, so assume that it is false. This means that $T_{\mathbf{H}_i}$ is true in \mathbf{P} for some i . Hence there exists a τ -homomorphism from \mathbf{H}_i to \mathbf{P} . It follows that we cannot have a τ -homomorphism from \mathbf{P} to \mathbf{Q} . Conversely, suppose that there is no homomorphism from \mathbf{P} to \mathbf{Q} and the unary relations in \mathbf{P} are pairwise disjoint; thus there exists some i such that $\mathbf{H}_i \preceq \mathbf{P}$. Thus the sentence $T_{\mathbf{H}_i}$ is true in \mathbf{P} .

Next we show that (1) holds. Let $\mathbf{H}_1, \dots, \mathbf{H}_l$ denote all the obstructions for \mathbf{Q} such that each $R_i^{\mathbf{H}_i}$ contains at most one element of H_i . Consider the sentence

$$\bigvee_{i=1}^l S_{\mathbf{H}_i}.$$

We claim that a poset \mathbf{P} containing \mathbf{Q} does not retract onto \mathbf{Q} if and only if this sentence is true in \mathbf{P} . Suppose that $S_{\mathbf{H}_i}$ is true in \mathbf{P} for some i . This implies that there is a τ -homomorphism from \mathbf{H}_i to \mathbf{P} (viewed as a τ -structure), and consequently there cannot be a τ -homomorphism from \mathbf{P} to \mathbf{Q} ; it follows that there is no σ -homomorphism either.

For the converse, suppose that \mathbf{P} does not retract onto \mathbf{Q} ; this means that there is no σ -homomorphism of \mathbf{P} to \mathbf{Q} , and no τ -homomorphism either. Now proceed exactly as in the proof of (2): there exists an obstruction \mathbf{H}_i of \mathbf{Q} such that $\mathbf{H}_i \preceq \mathbf{P}$; and simply notice that, since \mathbf{P} is a σ -structure, the obstruction \mathbf{H}_i will have at most one element in each unary relation. ■

5 Proof of Proposition 3.3

We shall begin with a general construction akin to the well-known attaching construction in topology (see, e.g., Definition 13.13 in [3]). Let \mathbf{T} and \mathbf{Q} be graphs, let $C \subseteq T$ be a subgraph of T and let $\phi : C \rightarrow Q$ be edge-preserving. We construct a new graph \mathbf{T}_ϕ obtained by glueing \mathbf{T} and \mathbf{Q} , identifying elements of C with their corresponding images in Q . More formally, let \mathbf{K} denote the disjoint union of the graphs \mathbf{T} and \mathbf{Q} , and define a partial function ϕ' from K to Q by

$$\phi'(x) = \begin{cases} x, & \text{if } x \in Q, \\ q, & \text{if } \phi(x) = q. \end{cases}$$

Let C' denote the domain of ϕ' . Notice that by definition $\phi'(x) = \phi(x)$ for all $x \in C$. It is clear that ϕ' is edge-preserving on C' . Define an equivalence relation on K as follows: let $x \sim y$ if $x = y$ or $\phi'(x) = \phi'(y)$. The base set of the graph \mathbf{T}_ϕ is the set of equivalence classes of the relation \sim ; denote the class of element x by $[x]$. We declare $[x][y]$ to be an edge of \mathbf{T}_ϕ if uv is an edge of \mathbf{K} for some $u \in [x]$ and some $v \in [y]$. (It is immediate that the resulting binary relation is reflexive and symmetric.)

Fact 5.1 (i) \mathbf{T}_ϕ contains a copy \mathbf{Q}' of \mathbf{Q} ;

(ii) \mathbf{T}_ϕ retracts onto \mathbf{Q}' if and only if there exists a homomorphism $\bar{\phi} : \mathbf{T} \rightarrow \mathbf{Q}$ that extends ϕ .

Proof. (i) Let $q_1, q_2 \in Q$. Then obviously $[q_1] = [q_2]$ if and only if $q_1 = q_2$. Now suppose that $[q_1][q_2]$ is an edge of \mathbf{T}_ϕ . This means there exist $q_1 \sim u$ and $v \sim q_2$ where uv is an edge of \mathbf{K} ; but then $q_1 = \phi'(u)$ and $\phi'(v) = q_2$ implies that $q_1 q_2$ is an edge of \mathbf{Q} . (ii) If r is a retraction of \mathbf{T}_ϕ onto \mathbf{Q}' define $\bar{\phi}(x) = i(r([x]))$ where i is the isomorphism from \mathbf{Q}' to \mathbf{Q} ; it is easy to verify that this is the desired map. Conversely, let $\bar{\phi}$ be an extension of ϕ . Clearly $\bar{\phi}$ is an extension of ϕ' when this last map is restricted to T . Define $r([x]) = [\bar{\phi}(x)]$ if $x \notin Q$ and $r([x]) = [x]$ otherwise. This is well defined: indeed, if $[x] = [y]$ and without loss of generality $x \notin Q$, then $\phi'(x) = \phi'(y)$ and clearly we have $\bar{\phi}(x) = \bar{\phi}(y)$ if $y \notin Q$; otherwise we certainly have $\phi'(x) = y = \bar{\phi}(x)$. Now we show that r is a homomorphism: let $[x][y]$ be an edge, i.e. $x \sim u$ and $v \sim y$ where uv is an edge of \mathbf{K} . But then $r([x]) = [\bar{\phi}(u)]$ and $[\bar{\phi}(v)] = r([y])$ and we are done. ■

5.1 The construction of the graphs

Let \mathbf{Q} be a graph which is either disconnected or does not admit an NU polymorphism. By Theorem 2.1 there exists an idempotent subalgebra X of a finite power of \mathbf{Q} which is not connected (if \mathbf{Q} is disconnected take $X = Q$). More precisely, there exists a triple $(\mathbf{Y}, (y_1, \dots, y_n), \gamma)$ where \mathbf{Y} is a graph, $y_i \in Y$ for all $i = 1, \dots, n$ and γ is a partial map from Y to Q with domain Y' with the following property: if X denotes the subset of Q^n that consists of all n -tuples

$$(\delta(y_1), \dots, \delta(y_n))$$

where $\delta : \mathbf{Y} \rightarrow \mathbf{Q}$ runs through homomorphisms whose restriction to Y' is equal to γ , then X is not connected. Let (x_1, \dots, x_n) and (x'_1, \dots, x'_n) be in distinct components of X .

Notice that by choosing n as small as possible we may assume that $y_i \notin Y'$ for all $1 \leq i \leq n$; for otherwise we could simply project X onto the remaining coordinates to obtain a disconnected idempotent subalgebra in a smaller power of \mathbf{Q} .

Let $l = 2p \geq 2$ be an even integer, and let \mathbf{C}_l denote the reflexive cycle on l elements, i.e. the graph on the set $\{0, 1, \dots, l-1\}$ where ij is an edge if and only if $|i-j| \leq 1$ (sum modulo l). Consider also the graph $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$, the disjoint union of two cycles on p elements. We shall assume that the underlying set of vertices of this graph is the same as that of \mathbf{C}_l , where $\{0, \dots, p-1\}$ will be one copy of \mathbf{C}_p and $\{p, \dots, 2p-1\}$ the other.

Lemma 5.2 *Let $m \geq 1$, $l = 2 \cdot 3^{m+3}$, and $p = 3^{m+3}$. Then the graphs \mathbf{C}_l and $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$ are m -isomorphic via a sequence $(I_j)_{j \leq m}$ such that $f(0) = 0$ and $f(p) = p$ for all $f \in I_j$ and all j .*

Proof: Let \mathbf{A} and \mathbf{A}' be obtained from \mathbf{C}_l and $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$, respectively, by endowing both of them with two unary relations $\{0\}$ and $\{p\}$. By Hanf's Locality Lemma (see [12] Theorem 2.4.1) structures \mathbf{A} and \mathbf{A}' are $(m+2)$ -isomorphic. To see this, notice that both graphs have exactly the same number of occurrences of every 3^{m+2} -ball type: one, if the ball is centered at 0 or p , two if it is not centered at 0 or p but contains one of them and finally $(2 \cdot 3^{m+3} - 4 \cdot 3^{m+2} - 2)$ if the ball does not contain 0 or p . Let $(I'_j)_{j \leq m+2}$ be the sequence whose existence is guaranteed by Hanf's Lemma. For any $0 \leq j \leq m$, let I_j contain those f in I'_j such that $f(0) = 0$ and $f(p) = p$. By applying the forth property on I'_{m+2} with 0 and in I'_{m+1} with p we can guarantee that I_m (and consequently $(I_j)_{j \leq m}$) is nonempty. ■

We shall now construct graphs \mathbf{S} and \mathbf{T} starting from the above graphs. We shall take n disjoint copies of the cycle \mathbf{C}_l (respectively the union of cycles $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$) and we glue l copies of \mathbf{Y} in the following manner

(see Fig. 1): the element y_i of the j -th copy of \mathbf{Y} is identified to the element j of the i -th cycle (respectively, of the i -th union of cycles $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$). More precisely, let \mathbf{U} be the disjoint union of l copies of \mathbf{Y} , say $\mathbf{Y} \times \{z\}$ for $z \in \{0, \dots, l-1\}$. Let \mathbf{C} (resp. \mathbf{D}) denote the disjoint union of n cycles \mathbf{C}_l (respectively the union of cycles $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$), say $\{t\} \times \mathbf{C}_l$ (resp. $\{t\} \times (\mathbf{C}_p \dot{\cup} \mathbf{C}_p)$) for $t \in \{1, \dots, n\}$. Let μ (resp. ν) be the partial map from U to C (respectively to D) that sends (y_i, z) to (i, z) ; then \mathbf{S} is the graph \mathbf{U}_μ , and \mathbf{T} is the graph \mathbf{U}_ν .

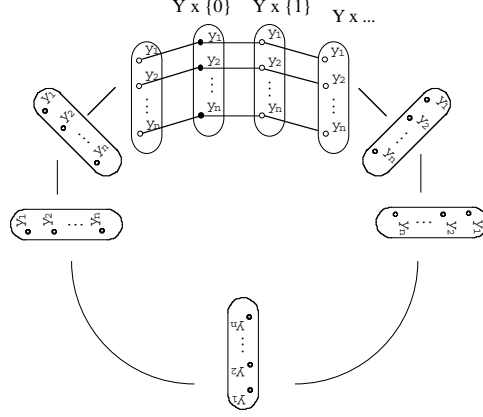


Figure 1: A partial view of the coloured graphs \mathbf{Y} glued to the n cycles \mathbf{C}_l .

Another way of viewing this: the graph \mathbf{S} (and similarly for \mathbf{T}) is obtained from the disjoint copies $\mathbf{Y} \times \{z\}$ by adding the edges $(y_i, j)(y_i, k)$ when jk is an edge in \mathbf{C}_l .

Define partial maps ϕ and ψ from \mathbf{S} and \mathbf{T} , respectively, to \mathbf{Q} as follows:

$$\phi(t) = \psi(t) = \begin{cases} \gamma(y), & \text{if } t = (y, z) \text{ for some } y \in Y', \\ x_i, & \text{if } t = (y_i, 0), \\ x'_i, & \text{if } t = (y_i, p). \end{cases}$$

Fact 5.3 *The maps ϕ and ψ are edge-preserving.*

Proof. We prove the result for ϕ . Let $t_1 t_2$ be an edge in the domain of ϕ . Notice that these elements must be in the same copy of \mathbf{Y} , i.e. $t_1 = (u, z)$ and $t_2 = (v, z)$ where uv is an edge of \mathbf{Y} . If $z \neq 0$ and $z \neq p$ then ϕ is equal to γ and we are done. If $z = 0$, let δ be an extension of γ such that $\delta(y_i) = x_i$ for all $1 \leq i \leq n$: it must exist, since $(x_1, \dots, x_n) \in X$. Clearly ϕ is a restriction of δ (on $Y \times \{0\}$) and hence is edge-preserving. The case $z = p$ is identical. ■

Fact 5.4 *(i) The map ψ admits an edge-preserving extension to \mathbf{T} ; (ii) the map ϕ admits no edge-preserving extension to \mathbf{S} .*

Proof. (i) Since (x_1, \dots, x_n) and (x'_1, \dots, x'_n) belong to X , there exist extensions δ and δ' of γ from \mathbf{Y} to \mathbf{Q} that map (y_1, \dots, y_n) to (x_1, \dots, x_n) and (x'_1, \dots, x'_n) respectively. Define an extension β of ψ by

$$\beta(t) = \begin{cases} \delta(y), & \text{if } t = (y, z) \text{ with } z < p, \\ \delta'(y), & \text{if } t = (y, z) \text{ with } z \geq p. \end{cases}$$

Since our graphs are reflexive, this is an edge-preserving extension of ψ .

(ii) Suppose that there exists an edge-preserving extension β of ϕ ; then for any $0 \leq z \leq l - 1$ the restriction of β to $Y \times \{z\}$ can be seen an extension of γ , and hence the tuple $(\beta(y_1, z), \dots, \beta(y_n, z))$ is in X . But then we have a path $(\beta(y_1, 0), \dots, \beta(y_n, 0)), (\beta(y_1, 1), \dots, \beta(y_n, 1)), \dots, (\beta(y_1, p), \dots, \beta(y_n, p))$ in X from (x_1, \dots, x_n) to (x'_1, \dots, x'_n) , a contradiction with the choice of these two tuples. ■

Corollary 5.5 *The graph \mathbf{T}_ψ retracts onto \mathbf{Q}' ; the graph \mathbf{S}_ϕ does not retract onto \mathbf{Q}' .*

Proof. Follows from Facts 5.1, 5.3 and 5.4. ■

5.2 Proof of m -isomorphism

To finish the proof of Proposition 3.3 we will show that \mathbf{T}_ψ and \mathbf{S}_ϕ are the desired graphs \mathbf{P} and \mathbf{P}' (if we identify \mathbf{Q} with \mathbf{Q}'). It remains to prove that the graphs \mathbf{S}_ϕ and \mathbf{T}_ψ are m -isomorphic. Let $(I_j)_{j \leq m}$ be the sequence whose existence is guaranteed by Lemma 5.2: we proceed to construct a sequence $(\tilde{I}_j)_{j \leq m}$. Fix $0 \leq j \leq m$. For any $f \in I_j$, define a partial map \tilde{f} from S_ϕ to T_ψ as follows: let B be any element of S_ϕ . If B contains no element in the domain of ϕ , then $B = \{(y, z)\}$ and we put $\tilde{f}(B) = [(y, f(z))]$ provided z is in the domain of f , otherwise we leave $\tilde{f}(B)$ undefined; if B contains an element in the domain of ϕ , then it contains a unique element $q \in Q$ and we define $\tilde{f}(B) = [q]$. Finally, define $\tilde{I}_j = \{\tilde{f} : f \in I_j\}$ for all $0 \leq j \leq m$.

Lemma 5.6 *The graphs \mathbf{S}_ϕ and \mathbf{T}_ψ are m -isomorphic via $(\tilde{I}_j)_{j \leq m}$; moreover every partial isomorphism in every \tilde{I}_j in the sequence fixes every element of Q' .*

Proof. Fix some $0 \leq j \leq m$ and some $f \in I_j$. It is clear that the function \tilde{f} is well-defined and that it fixes all elements in Q' , i.e. $\tilde{f}([q]) = [q]$ for all $q \in Q$.

Claim 1. Let $B \in S_\phi$. Then $B \cap Q \neq \emptyset$ if and only if $\tilde{f}(B) \cap Q \neq \emptyset$.

Proof of Claim 1. One direction is obvious by definition of \tilde{f} . Now suppose that $\tilde{f}(B) \cap Q \neq \emptyset$ and let $(y, z) \in B$. If (y, z) is in the domain of ϕ we are done. Otherwise $\tilde{f}(B) = [(y, f(z))]$ intersects Q so either y is in the domain of γ , and hence (y, z) is in the domain of ϕ , or else $y = y_i$ for some i and $f(z) \in \{0, p\}$; by the properties of f we get that $z \in \{0, p\}$ which again shows that (y, z) is in the domain of ϕ . In both cases, we conclude by definition of \mathbf{S}_ϕ that $B \cap Q \neq \emptyset$.

Claim 2. The map \tilde{f} is injective.

Proof of Claim 2. Let $\tilde{f}(B) = \tilde{f}(B')$. There are two cases: (i) if $\tilde{f}(B) = \tilde{f}(B')$ intersects Q then by the last claim both B and B' intersect Q ; if $q \in B \cap Q$ and $q' \in B' \cap Q$ then by definition of \tilde{f} we have that $[q] = \tilde{f}(B) = \tilde{f}(B') = [q']$ and hence $q = q'$ so $B = B'$. (ii) if $\tilde{f}(B) = \tilde{f}(B')$ does not intersect Q then we have by Claim 1 that $B = \{(y, z)\}$ and $B' = \{(y', z')\}$ for some $y, y' \in Y$ and $z, z' \in C_l$. By definition of \tilde{f} and the last claim we get that $(y, f(z)) = (y', f(z'))$ and, since f is injective, we conclude that $(y, z) = (y', z')$ so we are done.

Claim 3. The map \tilde{f} is edge-preserving.

Proof of Claim 3. Let BB' be an edge of \mathbf{S}_ϕ , i.e. there exist $u \in B$ and $v \in B'$ such that uv is an edge of \mathbf{S} ; we show that $\tilde{f}(B)\tilde{f}(B')$ is an edge in \mathbf{T}_ψ (assuming both are defined). The result is obvious if both B and B' meet Q . Suppose next that neither meets Q : then $B = \{(y, z)\}$ and $B' = \{(y', z')\}$ and either (i) $z = z'$, and yy' is an edge of \mathbf{Y} so the result follows easily, or (ii) $y = y' = y_i$ for some $1 \leq i \leq n$, zz' is an edge of \mathbf{C}_l and $z \notin \{0, p\}$; thus $f(z)f(z')$ is an edge and the result follows once again. By symmetry we may now suppose without loss of generality that $B = [q]$ and $B' = \{(y', z')\}$. Then (y, z) is adjacent to (y', z') for some (y, z) in the domain of ϕ . Then there are two cases: (i) $z = z'$, and yy' is an edge of \mathbf{Y} . Then $\psi((y, f(z))) = q$, $(y, f(z))$ and $(y', f(z))$ are adjacent so $\tilde{f}(B)\tilde{f}(B')$ is an edge in \mathbf{T}_ψ ; (ii) $y = y' = y_i$ for some i , $z \in \{0, p\}$ and zz' is an edge of \mathbf{C}_l . But f is edge-preserving so $(y_i, f(z))$ is adjacent to $(y_i, f(z'))$ in \mathbf{T} and, since f fixes 0 and p , we get that $\psi((y_i, f(z))) = q$, so we are done.

Claim 4. The inverse of the map \tilde{f} is edge-preserving.

Proof of Claim 4. Suppose that $\tilde{f}(B)\tilde{f}(B')$ is an edge in \mathbf{T}_ψ , i.e. there are elements $u \in \tilde{f}(B)$ and $v \in \tilde{f}(B')$ such that uv is an edge of \mathbf{T} . We show that BB' is an edge of \mathbf{S}_ϕ . Suppose first that both blocks meet Q , say $q \in \tilde{f}(B)$ and $q' \in \tilde{f}(B')$ where $q, q' \in Q$. Then of course qq' is an edge of \mathbf{Q} , and by Claim 1 and definition of \tilde{f} we get that $B = [q]$ is adjacent to $[q'] = B'$. Secondly, suppose that neither block meets Q . Then $\tilde{f}(B) = \{(y, f(z))\}$ and $\tilde{f}(B') = \{(y', f(z'))\}$ where $(y, f(z))(y', f(z'))$ is an edge in \mathbf{T} . This means that either (i) $f(z) = f(z')$ and yy' is an edge of \mathbf{Y} , and, since f is injective, we get that $z = z'$ so $(y, z)(y', z')$ is an edge of \mathbf{S} , or (ii) $y = y' = y_i$ for some i and $f(z)f(z')$ is an edge of $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$. Since $f \in I_j$, we get that zz' is an edge of \mathbf{C}_l and the rest follows easily. Finally suppose that $\tilde{f}(B) = [q]$ and $\tilde{f}(B') = \{(y', f(z'))\}$. Then there exists some (y, w) in the domain of ψ such that $\psi((y, w)) = q$ and $(y, w)(y', f(z'))$ is an edge in \mathbf{T} . If $w = f(z')$ and yy' is an edge of \mathbf{Y} , then by Claim 1 $(y', z') \in B'$ and it is adjacent to (y, z') in \mathbf{S} , and either y is in the domain of γ or $y = y_i$ and $z' = f(z') \in \{0, p\}$ so $(y, z') \in B$ and we are done. Otherwise we have that $y = y' = y_i$ for some i , $wf(z')$ is an edge of $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$ and $w \in \{0, p\}$. But then $f(w) = w$ and f is a partial isomorphism so wz' is an edge of \mathbf{C}_l , and the rest follows easily.

We have now proved that the sets \tilde{I}_j consist only of partial isomorphisms. Finally, we prove:

Claim 5. The sequence $(\tilde{I}_j)_{j \leq m}$ has the ‘back and forth’ property.

Proof of Claim 5. Let $j < m$ and let $\tilde{f} \in I_{j+1}$. Let $B \in S_\phi$ (the other case, $B' \in T_\psi$, is identical.) If B intersects Q then it is in the domain of \tilde{f} and we are done (we may certainly suppose that the sequence $(\tilde{I}_j)_{j \leq m}$ is decreasing.) So now assume that B does not meet Q , so $B = \{(y, z)\}$ where (y, z) is not in the domain of \tilde{f} , which means that z is not in the domain of f . By the back and forth property of the sequence $(I_j)_{j \leq m}$ we can find $g \in I_j$ such that z is in its domain and g is an extension of f . But then it is clear that (y, z) is in the domain of $\tilde{g} \in \tilde{I}_j$ and that \tilde{g} is an extension of \tilde{f} . ■

6 Proof of Proposition 4.2

6.1 A construction

Let \mathbf{T} and \mathbf{Q} be posets, let $C \subseteq T$ and let $\phi : C \rightarrow Q$ be a monotone map. We construct a new poset \mathbf{T}_ϕ obtained by glueing \mathbf{T} and \mathbf{Q} , identifying an element of C with its image in Q . More formally, let \mathbf{K} denote the disjoint union of the posets T and Q , and define a partial function ϕ' from K to Q by

$$\phi'(x) = \begin{cases} x, & \text{if } x \in Q, \\ q, & \text{if there exist } u \leq x \leq v \text{ such that} \\ & \phi(u) = \phi(v) = q. \end{cases}$$

Let C' denote the domain of ϕ' . Notice that by definition $\phi'(x) = \phi(x)$ for all $x \in C$.

A subset A of a poset \mathbf{P} is called *convex* if $b \in A$ whenever $a \leq b \leq c$ with $a, c \in A$.

Fact 6.1 (i) $\phi'^{-1}(q)$ is convex in \mathbf{K} for all $q \in Q$; (ii) ϕ' is monotone on C' .

Proof. (i) Let $x \leq y \leq z$ with $\phi'(x) = \phi'(z) = q$. If any of x, y, z is equal to q then so are the others so we are done. Otherwise, by definition of ϕ' there exist $u \leq x$ and $z \leq v$ such that $\phi(u) = \phi(v) = q$ hence $\phi'(y) = q$. (ii) Let $x \leq y$ in C' . If one of these is in Q we are done. Otherwise, there exist $u \leq x$ and $y \leq v$ such that $\phi'(x) = \phi(u) \leq \phi(v) = \phi'(y)$. ■

Define an equivalence relation on K as follows: let $x \sim y$ if $x = y$ or $\phi'(x) = \phi'(y)$. The base set of the poset \mathbf{T}_ϕ is the set of equivalence classes of the relation \sim ; denote the class of element x by $[x]$. We define $[x] \sqsubseteq [y]$ if $u \leq v$ for some $u \in [x]$ and some $v \in [y]$. The ordering on \mathbf{T}_ϕ is the transitive closure of the relation \sqsubseteq . (In what follows, we shall denote the ordering on any poset by the same symbol \leq , as usual.)

Fact 6.2 The relation \sqsubseteq defines a partial order.

Proof. The relation is clearly reflexive and transitive so we must show that \sqsubseteq is acyclic. Suppose that we have a sequence of elements of K as follows:

$$x_1 \sim y_1 \leq x_2 \sim y_2 \leq x_3 \cdots x_n \sim y_n \leq x_{n+1} \sim x_1.$$

We must show that $[x_1] = [x_2] = \cdots = [x_n]$. If $x_i = y_i$ for all i but one then we are done. Otherwise, there exists some $i < j$ such that $x_i \neq y_i$ and $x_j \neq y_j$. Choose $i < j$ as close to one another as possible in the cycle: then we get that $y_i \leq x_{i+1} \leq \cdots \leq x_{j-1} \leq x_j$ so that $\phi'(x_i) \leq \phi'(x_j)$ by Fact 1 (ii). Repeating this argument for all indices k such that $x_k \neq y_k$ shows that ϕ' is constant on those elements of the sequence where it is defined; and by Fact 1 (i) it follows that in fact ϕ' is defined for all elements of the sequence. ■

Fact 6.3 (i) \mathbf{T}_ϕ contains a copy \mathbf{Q}' of \mathbf{Q} ; (ii) \mathbf{T}_ϕ retracts onto \mathbf{Q} if and only if there exists a monotone map $\bar{\phi} : \mathbf{T} \rightarrow \mathbf{Q}$ that extends ϕ .

Proof. (i) Let $q_1, q_2 \in Q$. Then obviously $[q_1] = [q_2]$ if and only if $q_1 = q_2$. Now suppose that $[q_1] \sqsubseteq [q_2]$ in \mathbf{T}_ϕ . This means there exist $q_1 \sim u \leq v \sim q_2$. But then $q_1 = \phi'(u) \leq \phi'(v) = q_2$. (ii) If r is a retraction of \mathbf{T}_ϕ onto \mathbf{Q} define $\bar{\phi}(x) = r([x])$; it is easy to verify that this is the desired map. Conversely, let $\bar{\phi}$ be an extension of ϕ . Clearly $\bar{\phi}$ is an extension of ϕ' when this last map is restricted to T . Define $r([x]) = [\bar{\phi}(x)]$ if $x \notin Q$ and $r([x]) = [x]$ otherwise. This is well defined: indeed, if $[x] = [y]$ and without loss of generality $x \notin Q$, then $\phi'(x) = \phi'(y)$ and clearly we have $\bar{\phi}(x) = \bar{\phi}(y)$ if $y \notin Q$; otherwise we certainly have $\phi'(x) = y = \bar{\phi}(x)$. Now we show that r is monotone. Let $[x] \sqsubseteq [y]$, i.e. $x \sim u \leq v \sim y$. But then $r([x]) = [\bar{\phi}(u)] \sqsubseteq [\bar{\phi}(v)] = r([y])$. ■

6.2 The construction of the posets

Let \mathbf{Q} be a poset which is either disconnected or does not admit an NU polymorphism. By Theorem 2.1 there exists an idempotent subalgebra X of some power \mathbf{Q}^n which is not connected (if \mathbf{Q} is disconnected take $X = Q$). In fact, by a result of [24] we may assume that $n = 1$. More precisely, there exists a triple $(\mathbf{Y}, y_0, \gamma)$ where \mathbf{Y} is a poset, $y_0 \in Y$ and γ is a partial map from Y to Q with domain Y' with the following property: if X denotes the subset of Q that consists of all $\delta(y_0)$ where $\delta : \mathbf{Y} \rightarrow \mathbf{Q}$ runs through all monotone maps whose restriction to Y' is equal to γ , then X is not connected. Let x and x' be in distinct components of X .

We claim that we may choose (Y, y_0, γ) with the following properties: (i) if $y \in Y'$ is comparable to y_0 then $\gamma(y) \notin \{x, x'\}$ and (ii) if $y < y'$ in Y' then $\gamma(y) < \gamma(y')$.

Indeed, the first condition follows from the fact that x and x' are both in X and are incomparable. For the second statement: clearly γ is monotone on its domain (since X is nonempty) so if (ii) does not hold then we have $\gamma(y) = \gamma(y') = q$ for some $y < y'$ in Y' and some $q \in Q$. Obviously we may assume in that case that $\gamma(u) = q$ for all $y \leq u \leq y'$. It is easy to see that one may ‘fuse’ all these elements into one to obtain a new triple $(\mathbf{Y}_1, y_0, \gamma_1)$ with the same properties as before, namely that X is the set of all $\delta(y_0)$, where δ ranges over the set of all monotone extensions of γ_1 (simply define a partial map α from \mathbf{Y} to the one-element poset with domain $\{u : y \leq u \leq y'\}$ and use the construction of section 6.1).

Let p be an even integer and $l = 2p$. Let \mathbf{C}_l denote the crown on l elements, that is, the poset on the set $\{0, 1, \dots, l-1\}$ where $i < j$ if and only if i is even and $|i - j| = 1$ (sum modulo l). Consider also the graph $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$, the disjoint union of two crowns on p elements. We shall assume that the underlying set of vertices of this poset is the same as that of \mathbf{C}_l , where $\{0, \dots, p-1\}$ will be one copy of \mathbf{C}_p and $\{p, \dots, 2p-1\}$ the other (see Figure 2).




Figure 2: The posets \mathbf{C}_{16} and $\mathbf{C}_8 \dot{\cup} \mathbf{C}_8$.

We now view the above as coloured posets, i.e. structures with one binary relation (their ordering) and two constants, namely 0 and p . We claim that these structures are m -isomorphic.

Lemma 6.4 *Let $m \geq 1$, $l = 4 \cdot 3^{m+3}$, and $p = 2 \cdot 3^{m+3}$. Then the posets \mathbf{C}_l and $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$ are m -isomorphic via a sequence $(I_j)_{j \leq m}$ such that $f(0) = 0$ and $f(p) = p$ for all $f \in I_j$ and all j .*

Proof. The proof is almost identical to that of Lemma 5.2 above. The only difference is that l and p are doubled here, since the size of a crown is always an even number. ■

We shall now construct posets \mathbf{S} and \mathbf{T} starting from the above posets. We glue copies of \mathbf{Y} (at y_0) to every element of the crown and the union of crowns. More precisely, let \mathbf{U} be the disjoint union of l copies

of \mathbf{Y} , say $Y \times \{z\}$ for $z \in \{0, \dots, l-1\}$. Let μ (resp. ν) be the partial map from \mathbf{U} to \mathbf{C}_l (respectively to $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$) that sends (y_0, z) to z ; then \mathbf{S} is the poset \mathbf{U}_μ , and \mathbf{T} is the poset \mathbf{U}_ν .




Figure 3: A coloured poset \mathbf{Y} , and a partial view of the poset \mathbf{S} . Darkened vertices in the copies of \mathbf{Y} are in the domain of ψ .

Define partial maps ϕ and ψ from S and T , respectively, to Q , as follows:

$$\phi(t) = \psi(t) = \begin{cases} \gamma(y), & \text{if } t = (y, z) \text{ for some } y \in Y', \\ x, & \text{if } t = (y_0, 0), \\ x', & \text{if } t = (y_0, p). \end{cases}$$

Fact 6.5 *The maps ϕ and ψ are monotone.*

Proof. We prove the result for ϕ . Let $t_1 \leq t_2$ be in the domain of ϕ . Suppose first that these elements are in the same copy of Y , i.e. $t_1 = (y_1, z)$ and $t_2 = (y_2, z)$ where $y_1 \leq y_2$. If neither is equal to y_0 then $\phi(t_1) = \gamma(y_1) \leq \gamma(y_2) = \phi(t_2)$. Otherwise suppose without loss of generality that $y_1 = y_0$ and that $z = 0$; we must show that $x \leq \gamma(y_2)$. But, since $x \in X$, there exists an monotone extension $\bar{\gamma}$ of γ such that $\bar{\gamma}(y_0) = x$ so we are done.

Now suppose that t_1 and t_2 are in different copies of Y . Then we have that $t_1 = (y_1, z_1)$ and $t_2 = (y_2, z_2)$ where $y_1 \leq y_0 \leq y_2$ and $z_1 \leq z_2$. By definition of ϕ at most one of y_1, y_2 is equal to y_0 and hence the preceding argument applies here as well. ■

Fact 6.6 (i) *The map ψ admits a monotone extension to \mathbf{T} ; (ii) *The map ϕ admits no monotone extension to \mathbf{S} .**

Proof. (i) Since $x, x' \in X$, there exist extensions γ_x and $\gamma_{x'}$ of γ from \mathbf{Y} to \mathbf{Q} that map y_0 to x and x' respectively. Define an extension β of ψ by

$$\beta(t) = \begin{cases} \gamma_x(y), & \text{if } t = (y, z) \text{ with } z < p, \\ \gamma_{x'}(y), & \text{if } t = (y, z) \text{ with } z \geq p. \end{cases}$$

It is easy to verify that this is a monotone extension of ψ .

(ii) Since x and x' are in distinct components of X and \mathbf{C}_l is connected, clearly our claim will follow if we can prove that any extension of ϕ must map every (y_0, z) to X . And indeed, if β is an extension of ϕ then its restriction to $Y \times \{z\}$ is an extension of γ , so we are done. ■

Corollary 6.7 *The poset \mathbf{T}_ψ retracts onto \mathbf{Q} ; the poset \mathbf{S}_ϕ does not retract onto \mathbf{Q} .*

Proof. Follows from Facts 6.3, 6.5 and 6.6. ■

We require one last auxiliary result before prove the posets \mathbf{S}_ϕ and \mathbf{T}_ψ are m -isomorphic. We show that the blocks of the equivalence relation involved in the construction of the posets \mathbf{S}_ϕ and \mathbf{T}_ψ have a simple structure: the only blocks with more than one element are those that contain an element in the domain of ϕ (ψ).

Lemma 6.8 *Any element B of S_ϕ (respectively T_ψ) is of the following form: (i) $B = \{(y, z)\}$ for some y which is not in the domain of γ and $z \notin \{0, p\}$, or (ii) $B = \{(y, z)\}$ where $z \in \{0, p\}$ and $y \neq y_0$ or (iii) $B = \{(y_1, z_1), \dots, (y_n, z_n), q\}$ where $\{(y_1, z_1), \dots, (y_n, z_n)\} = \phi^{-1}(q)$ (respectively $\psi^{-1}(q)$).*

Proof. We consider only the case S_ϕ , the other is identical. Recall that the elements of S_ϕ are the blocks of the equivalence defined by $u \sim v$ if $u = v$ or $\phi'(u) = \phi'(v)$. Suppose that $\phi'(u) = \phi'(v)$ where u is in S ; by definition of ϕ' this means that there exist elements $(y, z) \leq u \leq (y', z')$ and $q \in Q$ such that $\phi((y, z)) = \phi((y', z')) = q$. We have that $y \leq y_0 \leq y'$ and $z \leq z'$. It follows easily from our claim on the triple $(\mathbf{Y}, y_0, \gamma)$ that $y = y'$ and $z = z'$, and hence u is in the domain of ϕ . The claim follows easily from this fact. ■

We are now in a position to prove that the posets \mathbf{S}_ϕ and \mathbf{T}_ψ are m -isomorphic. Let $(I_j)_{j \leq m}$ be the sequence whose existence is guaranteed by Lemma 6.4: we proceed to construct a sequence $(\tilde{I}_j)_{j \leq m}$. Fix $0 \leq j \leq m$. For any $f \in I_j$, define a partial map \tilde{f} from S_ϕ to T_ψ as follows: by Lemma 6.8, the elements of S_ϕ are of two kinds: (a) if B contains no element in the domain of ϕ , then $B = \{(y, z)\}$ and we put $\tilde{f}(B) = [(y, f(z))]$ provided z is in the domain of f , otherwise we leave $\tilde{f}(B)$ undefined; (b) if B contains an element in the domain of ϕ , then it contains a unique element $q \in Q$ and we define $\tilde{f}(B) = [q]$. Finally, define $\tilde{I}_j = \{\tilde{f} : f \in I_j\}$ for all $0 \leq j \leq m$.

Lemma 6.9 *The posets \mathbf{S}_ϕ and \mathbf{T}_ψ are m -isomorphic via $(\tilde{I}_j)_{j \leq m}$; moreover every partial isomorphism in every I_j of the sequence fixes every element of Q .*

Proof. Fix some $0 \leq j \leq m$ and some $f \in I_j$. It is clear that the function \tilde{f} is well-defined and that it fixes all elements in Q , i.e. $\tilde{f}([q]) = [q]$ for all $q \in Q$.

Claim 1. Let $B \in S_\phi$. Then $B \cap Q \neq \emptyset$ if and only if $\tilde{f}(B) \cap Q \neq \emptyset$.

Proof of Claim 1. One direction is obvious by definition of \tilde{f} . Now suppose that $\tilde{f}(B) \cap Q \neq \emptyset$ and let $(y, z) \in B$. By Lemma 6.8 $(y, f(z))$ is in the domain of ψ , i.e. either y is in the domain of γ , and hence (y, z) is in the domain of ϕ , or else $y = y_0$ and $f(z) \in \{0, p\}$; by the properties of f we get that $z \in \{0, p\}$ which again shows that (y, z) is in the domain of ϕ . In both cases, we conclude by definition of S_ϕ that $B \cap Q \neq \emptyset$.

Claim 2. The map \tilde{f} is injective.

Proof of Claim 2. Let $\tilde{f}(B) = \tilde{f}(B')$. By Lemma 6.8 there are two cases: (i) if $\tilde{f}(B) = \tilde{f}(B')$ intersects Q then by the last claim both B and B' intersect Q ; if $q \in B \cap Q$ and $q' \in B' \cap Q$ then by definition of \tilde{f} we have that $[q] = \tilde{f}(B) = \tilde{f}(B') = [q']$ and hence $q = q'$ so $B = B'$. (ii) if $\tilde{f}(B) = \tilde{f}(B')$ do not intersect Q then we have by Claim 1 that $B = \{(y, z)\}$ and $B' = \{(y', z')\}$ for some $y, y' \in Y$ and $z, z' \in C_l$. By definition of \tilde{f} and the last claim we get that $(y, f(z)) = (y', f(z'))$ and, since f is injective, we conclude that $(y, z) = (y', z')$ so we are done.

Claim 3. The map \tilde{f} is monotone.

Proof of Claim 3. Let $B \leq B'$ in \mathbf{S}_ϕ , i.e. let there be a sequence of blocks $B = B_1, \dots, B_n = B'$ and elements $u_i, v_i \in B_i$ such that

$$u_1 \leq v_2 \sim u_2 \leq v_3 \cdots v_{n-1} \sim u_{n-1} \leq v_n.$$

Suppose that the block B_i is not in the domain of \tilde{f} . This implies that it cannot meet Q and hence by Lemma 6.8 we have that

$$u_{i-1} \leq v_i = u_i \leq v_{i+1}$$

so that block B_i may actually be removed from the sequence. Hence it will suffice to prove that if B and B' are such that there exist $u \in B$ and $v \in B'$ such that $u \leq v$ then $\tilde{f}(B) \leq \tilde{f}(B')$ in \mathbf{T}_ψ .

The result is obvious if both B and B' meet Q . Suppose now that neither meets Q : then $B = \{(y, z)\}$ and $B' = \{(y', z')\}$ where $y \leq y_0 \leq y'$ and $z \leq z'$; thus $f(z) \leq f(z')$ and the result follows easily. Suppose now that $B = [q]$ and $B' = \{(y', z')\}$. Then $(y, z) \leq (y', z')$ for some (y, z) in the domain of ϕ . This means that $y \leq y_0 \leq y'$ and $z \leq z'$. If y is in the domain of γ , then $(y, f(z')) \leq (y', f(z'))$ in \mathbf{T}_ϕ , and $(y, f(z'))$ is in the domain of ψ ; by definition of ϕ and ψ we get that

$$\psi((y, f(z'))) = \gamma(y) = \phi((y, z)) = q$$

hence

$$\tilde{f}(B) = [q] = [(y, f(z'))] \leq [(y', f(z'))] = \tilde{f}(B').$$

Otherwise, we have that $y = y_0$ and $z \in \{0, p\}$; we suppose that $z = 0$ the other case being identical. By Lemma 6.4 we have that $0 = f(0) \leq f(z')$ so $(y, z) \leq (y', f(z'))$ and thus $\tilde{f}(B) \leq \tilde{f}(B')$ in \mathbf{T}_ψ . The last case, where $B = \{(y, z)\}$ and $B' = [q]$, is quite similar.

Claim 4. The inverse of the map \tilde{f} is monotone.

Proof of Claim 4. Suppose that $\tilde{f}(B) \leq \tilde{f}(B')$ in \mathbf{T}_ψ , i.e. let there be a sequence of blocks $\tilde{f}(B) = B_1, \dots, B_n = \tilde{f}(B')$ and elements $u_i, v_i \in B_i$ such that

$$u_1 \leq v_2 \sim u_2 \leq v_3 \cdots v_{n-1} \sim u_{n-1} \leq v_n.$$

As in Claim 3, it is easy to see that we may assume that every block B_i is in the image of \tilde{f} ; and hence it will suffice to prove the following assertion: if there exist elements $u \in \tilde{f}(B)$ and $v \in \tilde{f}(B')$ such that $u \leq v$ then $B \leq B'$ in \mathbf{S}_ϕ . Suppose first that both blocks meet Q , say $q \in \tilde{f}(B)$ and $q' \in \tilde{f}(B')$ where $q, q' \in Q$. Then of course $q \leq q'$, and by Claim 1 and definition of \tilde{f} we get that $B = [q] \leq [q'] = B'$. Secondly, suppose that neither block meets Q . Then $\tilde{f}(B) = \{(y, f(z))\}$ and $\tilde{f}(B') = \{(y', f(z'))\}$ where $(y, f(z)) \leq (y', f(z'))$ in \mathbf{T} . This means that $y \leq y_0 \leq y'$ in \mathbf{Y} and $f(z) \leq f(z')$ in $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$. Since $f \in I_j$, we get that $z \leq z'$ and the rest follows easily. Thirdly suppose that $\tilde{f}(B) = [q]$ and $\tilde{f}(B') = \{(y', f(z'))\}$. Then there exists some (y, w) in the domain of ψ such that $\psi((y, w)) = q$ and $(y, w) \leq (y', f(z'))$ in \mathbf{T} . This means that $y \leq y_0 \leq y'$ in \mathbf{Y} and $w \leq f(z')$ in $\mathbf{C}_p \dot{\cup} \mathbf{C}_p$. Suppose first that $\gamma(y) = q$. But then $(y, z') \leq (y', z')$ in \mathbf{S} and, since $\gamma(y) = q$, we obtain that

$$B = [q] = [(y, z')] \leq [(y', z')] = B'.$$

If on the other hand $y = y_0$ and $w \in \{0, p\}$, we get that (assuming once again without loss of generality that $w = 0$) $f(0) = 0 \leq f(z')$ and hence $0 \leq z'$; it follows that $(y_0, 0) \leq (y', z')$ in \mathbf{S} and thus

$$B = [q] = [(y_0, 0)] \leq [(y', z')] = B'.$$

The fourth case, where $\tilde{B} = \{(y, z)\}$ and $\tilde{B}' = [q]$ is similar.

We have now proved that the sets \tilde{I}_j consist only of partial isomorphisms. Finally, we prove:

Claim 5. The sequence $(\tilde{I}_j)_{j \leq m}$ has the ‘back and forth’ property.

Proof of Claim 5. Let $j < m$ and let $\tilde{f} \in I_{j+1}$. Let $B \in S_\phi$ (the other case, $B' \in T_\psi$, is identical.) If B intersects Q then it is in the domain of \tilde{f} and we are done (we may certainly suppose that the sequence $(\tilde{I}_j)_{j \leq m}$ is decreasing.) So now assume that B does not meet Q , so $B = \{(y, z)\}$ where (y, z) is not in the domain of \tilde{f} , which means that z is not in the domain of f . By the back and forth property of the sequence $(I_j)_{j \leq m}$ we may find $g \in I_j$ such that z is in its domain and g is an extension of f . But then it is clear that (y, z) is in the domain of $\tilde{g} \in \tilde{I}_j$ and that \tilde{g} is an extension of \tilde{f} . ■

7 Conclusion

We have completely characterised posets and reflexive graphs for which poset retraction and graph retraction problems, respectively, are definable in first-order logic. We believe that this line of research can be successfully continued by considering other logics and other classes of structures. The key to our results is Theorem 2.1 relating finiteness of obstructions and certain connectedness properties, which in the case of posets and reflexive graphs happens to be captured by NU polymorphisms. To make further progress in looking for an algebraic description of homomorphism and retraction problems in \mathbf{FO} , it seems necessary to obtain more information about how the two above properties are linked for more general structures. Results of [1, 29] may provide some insight into this.

Acknowledgments

We are thankful to Christelle Vincent for helpful discussions. Victor Dalmau was partially supported by MCyT grants TIC2002-04470-C03 and TIC2002-04019-C03. Andrei Krokhin was partially supported by UK EPSRC grant GR/R29598. Benoit Larose is partially funded by NSERC.

References

- [1] A. Atserias. On digraph coloring problems and treewidth duality. In *Proc. IEEE 20th Symp. on Logic in Computer Science, LICS'05*, pages 106–115, 2005.
- [2] M. Benke. Some complexity bounds for subtype inequalities. *Theoretical Computer Science*, 212:3–27, 1999.
- [3] G. Bredon. *Topology and Geometry*. Springer-Verlag, third edition, 1995.
- [4] R. Brewster, T. Feder, P. Hell, J. Huang, and G. MacGillavray. Near-unanimity functions and varieties of graphs. submitted, 2004.

- [5] A. Bulatov. Tractable conservative constraint satisfaction problems. In *Proc. 18th IEEE Symp. on Logic in Computer Science, LICS'03*, pages 321–330, 2003.
- [6] A. Chandra and P. Merlin. Optimal implementation of conjunctive queries in relational databases. In *Proc. 9th ACM Symp. on Theory of Computing, STOC'77*, pages 77–90, 1977.
- [7] N. Creignou, S. Khanna, and M. Sudan. *Complexity Classifications of Boolean Constraint Satisfaction Problems*, volume 7 of *SIAM Monographs on Discrete Mathematics and Applications*. 2001.
- [8] V. Dalmau. Linear Datalog and bounded path duality of relational structures. *Logical Methods in Computer Science* (electronic), 1, 2005.
- [9] V. Dalmau, Ph. Kolaitis, and M. Vardi. Constraint satisfaction, bounded treewidth, and finite-variable logics. In *Proc. 8th Int. Conf. on Constraint Programming, CP'02*, volume 2470 of *LNCS*, pages 310–326, 2002.
- [10] J. Demetrovics, L. Hannák, and L. Rónyai. Near-unanimity functions of partial orders. In *Proc. 14th Int. Symp. on Multiple-Valued Logic, ISMVL'84*, pages 52–56, 1984.
- [11] D. Duffus and I. Rival. A structure theory for ordered sets. *Discrete Mathematics*, 35:53–118, 1981.
- [12] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer-Verlag, 1999.
- [13] T. Feder and P. Hell. List homomorphisms to reflexive graphs. *Journal of Combinatorial Theory, Ser.B*, 72:236–250, 1998.
- [14] T. Feder and M. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28:57–104, 1998.
- [15] T. Feder and M. Vardi. Homomorphism closed vs. existential positive. In *Proc. 18th IEEE Symp. on Logic in Computer Science, LICS'03*, pages 311–320, 2003.
- [16] P. Hell. Algorithmic aspects of graph homomorphisms. In C. Wensley, editor, *Surveys in Combinatorics 2003*, volume 307 of *LMS Lecture Note Series*, pages 239 – 276. Cambridge University Press, 2003.
- [17] P. Hell and J. Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.
- [18] M. Jawhari, D. Misane, and M. Pouzet. Retracts: graphs and ordered sets from metric point of view. *Contemporary Mathematics*, 57:175–226, 1986.
- [19] P. Jeavons, D. Cohen, and M. Cooper. Constraints, consistency and closure. *Artificial Intelligence*, 101(1-2):251–265, 1998.
- [20] Ph. Kolaitis. On the expressive power of logics on finite models. In *Finite Model Theory and its Applications*, EATCS Series: Texts in Theoretical Computer Science. Springer, 2005. to appear.
- [21] Ph. Kolaitis and M. Vardi. On the expressive power of Datalog: tools and a case study. *Journal of Computer and System Sciences*, 51:110–134, 1995.
- [22] Ph. Kolaitis and M. Vardi. Conjunctive-query containment and constraint satisfaction. *Journal of Computer and System Sciences*, 61:302–332, 2000.
- [23] A. Krokhin and B. Larose. Solving order constraints in logarithmic space. In *Proc. 20th Int. Symp. on Theoretical Aspects of Computer Science, STACS'03*, volume 2607 of *LNCS*, pages 379–390, 2003.

- [24] G. Kun and C. Szabó. Order varieties and monotone retractions of finite posets. *Order*, 18:79–88, 2001.
- [25] B. Larose, C. Loten, and L. Zádori. A polynomial-time algorithm to recognise near-unanimity graphs. *J. Algorithms*, 55:177–191, 2005.
- [26] B. Larose and L. Zádori. Algebraic properties and dismantlability of finite posets. *Discrete Mathematics*, 163:89–99, 1997.
- [27] B. Larose and L. Zádori. The complexity of the extendibility problem for finite posets. *SIAM Journal on Discrete Mathematics*, 17(1):114–121, 2003.
- [28] P. Lincoln and J. Mitchell. Algorithmic aspects of type interference with subtypes. In *Proc. 19th ACM Symp. on Principles of Programming Languages, POPL'92*, pages 293–304, 1992.
- [29] J. Nešetřil and C. Tardif. Duality theorems for finite structures (characterising gaps and good characterisations). *Journal of Combinatorial Theory, Ser.B*, 80:80–97, 2000.
- [30] R. Pöschel and L. Kalužnin. *Funktionen- und Relationenalgebren*. DVW, Berlin, 1979.
- [31] V. Pratt and J. Tiuryn. Satisfiability of inequalities in a poset. *Fundamenta Informaticae*, 28:165–182, 1996.
- [32] B. Rossman. Existential positive types and preservation under homomorphisms. In *Proc. IEEE 20th Symp. on Logic in Computer Science, LICS'05*, pages 467–476, 2005.
- [33] N. Vikas. Compaction, retraction, and constraint satisfaction. *SIAM Journal on Computing*, 33:761–782, 2004.
- [34] L. Zádori. Relational sets and categorical equivalence of algebras. *Int. Journal of Algebra and Computation*, 7:561–576, 1997.