

Majority constraints have bounded pathwidth duality

Víctor Dalmau¹ and Andrei Krokhin²

¹ Department of Technology, University Pompeu Fabra, Barcelona 08003, Spain
`victor.dalmau@upf.edu`

² Department of Computer Science, University of Durham, Durham, DH1 3LE, UK
`andrei.krokhin@durham.ac.uk`

Abstract. We study certain constraint satisfaction problems which are the problems of deciding whether there exists a homomorphism from a given relational structure to a fixed structure with a majority polymorphism. We show that such a problem is equivalent to deciding whether the given structure admits a homomorphism from an obstruction belonging to a certain class of structures of bounded pathwidth. This implies that the constraint satisfaction problem for any fixed structure with a majority polymorphism is in **NL**.

Keywords: constraint satisfaction problem, homomorphism, complexity, pathwidth, bounded pathwidth duality, majority operation

1 Introduction and Related Work

The constraint satisfaction problem (CSP) provides a framework in which it is possible to express, in a natural way, many combinatorial problems encountered in artificial intelligence and computer science. A constraint satisfaction problem is represented by a set of variables, a domain of values for each variable, and a set of constraints between variables. The aim in a constraint satisfaction problem is then to find an assignment of values to the variables that satisfies the constraints.

It has been observed [10] (see also [14]) that the constraint satisfaction problem can be recast as the following fundamental problem: given two finite relational structures \mathbf{A} and \mathbf{B} , is there a homomorphism from \mathbf{A} to \mathbf{B} ? The CSP is **NP**-complete in general, and the identifying of its subproblems that have lower complexity has been a very active research direction in the last decade (see, e.g., [2, 4–6, 10, 11, 14, 19]).

One of the most studied restrictions on the CSP is when the structure \mathbf{B} is fixed, and only \mathbf{A} is part of the input. The obtained problem is denoted by $\text{CSP}(\mathbf{B})$. Examples of such problems include k -SAT, GRAPH H -COLOURING, and SYSTEMS OF EQUATIONS (e.g., linear equations). Strong motivation for studying this framework was given in [10] where it was shown that such problems can be

used in attempts to identify a largest subclass of \mathbf{NP} that avoids problems of intermediate complexity.

A variety of mathematical approaches to study problems $\text{CSP}(\mathbf{B})$ has been recently suggested. The most advanced approaches use logic (e.g., [18]), combinatorics (e.g., [11, 12]), universal algebra (e.g., [4, 19]), or combinations of those (e.g., [6, 8, 10, 21]).

Duality. The concept of *duality* has been much used to study homomorphism problems. The idea is to provide a set $\mathcal{O}_{\mathbf{B}}$ of *obstructions* for \mathbf{B} such that, for any relational structure \mathbf{A} , \mathbf{A} homomorphically maps to \mathbf{B} if and only if \mathbf{A} does not admit a homomorphism from any structure from $\mathcal{O}_{\mathbf{B}}$. If the set $\mathcal{O}_{\mathbf{B}}$ can be chosen so that it has certain nice properties, then the complexity of $\text{CSP}(\mathbf{B})$ is low. The forms of duality that have been considered in the literature include finite duality, bounded pathwidth duality, and bounded treewidth duality.

A structure \mathbf{B} has *finite duality* if there is a finite obstruction set $\mathcal{O}_{\mathbf{B}}$. Such dualities have been studied in [1, 21, 23]. The problems $\text{CSP}(\mathbf{B})$ with finite duality are exactly those for which the class of ‘yes’-instances is definable in first-order logic [1]. Clearly, such problems belong to the complexity class \mathbf{AC}^0 . A combinatorial characterisation of structures with finite duality is given in [23], and a universal-algebraic characterisation of such structures was obtained in [21].

Bounded pathwidth duality was introduced in [6, 7]³. A structure \mathbf{B} has bounded pathwidth duality if one can choose an obstruction set $\mathcal{O}_{\mathbf{B}}$ consisting of structures of bounded pathwidth (see the formal definition in Section 2). Several equivalent conditions, such as definability in various logics (e.g., in linear Datalog), and a useful connection of bounded pathwidth duality with certain games, called pebble-relation games, are given in [6, 7]. We will define and use this connection in Section 4. The problems $\text{CSP}(\mathbf{B})$ with bounded pathwidth duality belong to the complexity class \mathbf{NL} [6, 7]; moreover, all problems $\text{CSP}(\mathbf{B})$ known to belong to \mathbf{NL} have bounded pathwidth duality. Concrete examples of such problems are given in the two papers mentioned above, we will discuss some of them later. To the best of our knowledge, no characterisation of structures with bounded pathwidth duality is known.

Bounded treewidth duality has been studied in [3, 10, 12, 13, 18, 22]. This notion is similar to bounded pathwidth duality (and actually preceded and inspired it), but with $\mathcal{O}_{\mathbf{B}}$ consisting of structures of bounded treewidth. Many equivalent logical characterisations of bounded treewidth duality are known [10, 18], e.g., definability of the complement of a problem in Datalog. The problems $\text{CSP}(\mathbf{B})$ with bounded treewidth duality belong to the complexity class \mathbf{PTIME} [10]. No characterisation of structures with bounded treewidth duality is known, though there exists a strong necessary universal-algebraic condition, which is conjectured to also be sufficient [3, 22]. Clearly, finite duality implies bounded pathwidth duality, which, in turn, implies bounded treewidth duality.

Thus, obtaining necessary and/or sufficient algebraic conditions characterising structures with a given type of duality is a natural and interesting problem.

³ In [6, 7], it was called *bounded path duality*.

This paper contributes to the study of the problem for bounded pathwidth duality.

Algebraic approach. The algebraic approach to constraint satisfaction [2–4, 14, 19] is probably the most successful one. The key concept in this approach is the concept of a *polymorphism* (see the formal definition in Section 2) of a relational structure. A polymorphism is an operation which preserves each relation in the structure in the sense that it is a homomorphism from a finite Cartesian power of the structure to the structure itself. Two structures with the same polymorphisms have essentially the same properties with regard to the corresponding constraint satisfaction problems [4, 14, 19]. In particular, the problems have the same complexity, which makes polymorphisms very useful in classifying relational structures.

The existence of several forms of polymorphisms has been shown to guarantee that the corresponding CSPs are in **PTIME** (see [4, 14, 19]). One particular form of such polymorphism is a *majority* operation, which is a ternary operation ϕ on a set B satisfying $\phi(x, x, y) = \phi(x, y, x) = \phi(y, x, x) = x$ for all $x, y \in B$. Such operations have played an important role in earlier investigations. For example, the well-known LIST H -COLOURING problem for graphs [12] can be viewed as $\text{CSP}(\mathbf{B})$ for the structure \mathbf{B} whose relations are the binary edge relation of H and all possible unary relations. It is known that this problem $\text{CSP}(\mathbf{B})$ is in **PTIME** if and only if \mathbf{B} has a majority polymorphism [9]. As another example, consider the smallest non-trivial case, $|B| = 2$. If a two-element structure \mathbf{B} has a (unique) majority polymorphism, then $\text{CSP}(\mathbf{B})$ is a subproblem of the 2-SAT problem [16], and hence belongs to **NL**. So, the class of problems $\text{CSP}(\mathbf{B})$ with a structure \mathbf{B} having a majority polymorphism is a wide generalisation of 2-SAT.

Majority polymorphisms are known to guarantee bounded treewidth duality for relational structures (see [10] for the general situation, or Section 5.5 of [12] for graph H -colouring problems). Moreover, the corresponding CSPs are known to be solvable by a special sort of greedy algorithm (this property is referred to as ‘bounded strict width’ in [10]).

There exist structures with bounded pathwidth duality, but without a majority polymorphism [6, 8, 20], but many structures that are known to have bounded pathwidth duality, also have a majority polymorphism [6, 7]. For example, all oriented paths and directed cycles, have bounded pathwidth duality and also a majority polymorphism [6]. Implicational (or 0/1/all) constraints, introduced in [15], have a very particular form of majority polymorphism (called dual discriminator), and they have been shown to have bounded pathwidth duality [6, 7]. It was also shown in [6, 7] that a mild generalisation of the dual discriminator polymorphism also guarantees bounded pathwidth duality. Our main result shows, for the first time, that any structure with a majority polymorphism has bounded pathwidth duality, and hence the corresponding CSP belongs to **NL**. This answers an open question posed in [6, 7].

2 Basic Definitions

Most of the terminology introduced in this section is fairly standard. A *vocabulary* is a finite set of relation symbols or predicates. In what follows, τ always denotes a vocabulary. Every relation symbol R in τ has an *arity* $r = \rho(R) \geq 0$ associated to it. We also say that R is an r -ary relation symbol.

A τ -structure \mathbf{A} consists of a set A , called the *universe* of \mathbf{A} , and a relation $R^{\mathbf{A}} \subseteq A^r$ for every relation symbol $R \in \tau$ where r is the arity of R . All structures in this paper are assumed to be *finite*, i.e., structures with a finite universe. Throughout the paper we use the same boldface and slanted capital letters to denote a structure and its universe, respectively.

Let \mathbf{A} and \mathbf{A}' be τ -structures. We say that \mathbf{A}' is a substructure of \mathbf{A} , denoted by $\mathbf{A}' \subseteq \mathbf{A}$, if $A' \subseteq A$ and for every $R \in \tau$, $R^{\mathbf{A}'} \subseteq R^{\mathbf{A}}$. If \mathbf{A} is a τ -structure and $I \subseteq A$, then $\mathbf{A}|_I$ denotes the substructure induced by \mathbf{A} on I , i.e., the τ -structure \mathbf{I} with universe I and $R^{\mathbf{I}} = R^{\mathbf{A}} \cap I^r$ for every r -ary $R \in \tau$.

A *homomorphism* from a τ -structure \mathbf{A} to a τ -structure \mathbf{B} is a mapping $h : A \rightarrow B$ such that for every r -ary $R \in \tau$ and every $(a_1, \dots, a_r) \in R^{\mathbf{A}}$, we have $(h(a_1), \dots, h(a_r)) \in R^{\mathbf{B}}$. We denote this by $h : \mathbf{A} \rightarrow \mathbf{B}$. We say that \mathbf{A} homomorphically maps to \mathbf{B} , and denote this by $\mathbf{A} \rightarrow \mathbf{B}$ if there exists a homomorphism from \mathbf{A} to \mathbf{B} . Let a_1, \dots, a_m be elements in A and let b_1, \dots, b_m be elements in B . We shall write $\mathbf{A}, a_1, \dots, a_m \rightarrow \mathbf{B}, b_1, \dots, b_m$ to denote that there exists some homomorphism h from \mathbf{A} to \mathbf{B} such that $h(a_i) = b_i$, $1 \leq i \leq m$. For any subset I of A , any homomorphism from $\mathbf{A}|_I$ to \mathbf{B} is called a *partial homomorphism* from \mathbf{A} to \mathbf{B} .

Finally, $\text{CSP}(\mathbf{B})$ is defined to be the set of all structures \mathbf{A} such that $\mathbf{A} \rightarrow \mathbf{B}$.

Definition 1. A τ -structure \mathbf{A} is said to have treewidth at most (j, k) if there is a tree T , called a tree-decomposition of \mathbf{A} , such that

1. the nodes of T are subsets of A of size at most k ,
2. adjacent nodes can share at most j elements,
3. nodes containing any given element of A form a subtree,
4. for any tuple in any relation in \mathbf{A} , there is a node in T containing all elements from that tuple.

If T is a path then it is called a path-decomposition of \mathbf{A} , and \mathbf{A} is said to have pathwidth at most (j, k) .

Note that, traditionally, a structure is said to have pathwidth (treewidth) at most $k - 1$ if it has pathwidth (treewidth, respectively) at most (j, k) for some j , according to the above definition. Note that we use two numbers to parameterize treewidth and pathwidth, as is customary in the study of CSPs [10, 22] (rather than one as is customary in graph theory), for the following reason. The first parameter gives a more convenient parameterization of CSPs because the second parameter is bounded from below by the maximum arity of a relation in a structure, and hence it is less convenient to use for uniform treatment of structures of different signatures that behave essentially in the same way with

respect to homomorphisms. Nevertheless, the notions of pathwidth and treewidth of relational structures are closely related to the corresponding notions for graphs (see, e.g., [7]).

Definition 2. A set \mathcal{O} of τ -structures is called an obstruction set for \mathbf{B} if, for any τ -structure \mathbf{A} , $\mathbf{A} \rightarrow \mathbf{B}$ if and only if $\mathbf{A}' \not\rightarrow \mathbf{A}$ for all $\mathbf{A}' \in \mathcal{O}$.

A structure \mathbf{B} is said to have (j, k) -pathwidth duality if it has an obstruction set consisting of structures of pathwidth at most (j, k) . We say that \mathbf{B} has j -pathwidth duality if it has (j, k) -pathwidth duality for some $k > j$. We say that \mathbf{B} has bounded pathwidth duality if it has j -pathwidth duality for some $j \geq 0$.

By replacing “pathwidth” with “treewidth” throughout the above definition, one obtains the corresponding definitions of treewidth dualities.

Let us now formally define polymorphisms of relations and structures.

Definition 3. Let f be an n -ary operation on B , and R an m -ary relation on B . Then f is said to be a polymorphism of R (or R is invariant under f) if the following holds: for any $m \times n$ matrix X over B whose columns belong to R , the m -tuple $f(X)$ computed by applying f to the rows of X also belongs to R .

An operation is called a polymorphism of a relational structure if it is a polymorphism of every relation in the structure.

For example, it is well known, and easy to check, that any binary Boolean relation is invariant under the unique majority operation on $\{0, 1\}$.

One can easily check that f is an n -ary polymorphism of a relation structure \mathbf{B} if and only if $f : \mathbf{B} \times \dots \times \mathbf{B} \rightarrow \mathbf{B}$, where the product contains n copies of \mathbf{B} . If f is a polymorphism of a τ -structure \mathbf{B} , then one can generate from \mathbf{B} other relations that are invariant under f , as follows.

Lemma 1. Let \mathbf{B} be a τ -structure with a polymorphism f . Let \mathbf{C} be an arbitrary τ -structure, and fix arbitrary (not necessarily distinct) elements c_1, \dots, c_n of \mathbf{C} . Then the relation

$$\{(b_1, \dots, b_n) \mid \mathbf{C}, c_1, \dots, c_n \rightarrow \mathbf{B}, b_1, \dots, b_n\}$$

is also invariant under f .

The proof of this lemma is straightforward. (The construction in the above lemma is very similar to the indicator construction used in the study of H -colouring problems [12]).

Definition 4. An m -ary operation f on B has j -pathwidth duality if every structure with polymorphism f has j -pathwidth duality.

3 Main Result

Recall that a majority operation on a set B is a ternary operation ϕ on B which satisfies the identities $\phi(x, x, y) = \phi(x, y, x) = \phi(y, x, x) = x$ for all $x, y \in B$.

We will call a subset U of B a ϕ -subalgebra (or simply subalgebra if ϕ is clear from the context) if it is invariant under ϕ , that is, $\phi(x, y, z) \in U$ for all $x, y, z \in U$. Let U be a subalgebra of B , and $I \subseteq U$. Then we say that I is an ϕ -ideal (or, simply, an ideal) in U if $\phi(x, y, z) \in I$ provided $x, y, z \in U$, and at least two of them belong to I . For example, every subalgebra is an ideal in itself, and every singleton is an ideal in any subalgebra that contains it.

Relations invariant under a majority operation have the nice property of 2-decomposability. For an n -ary relation R on B , and for $1 \leq i < j \leq n$, let $\text{pr}_{i,j} R = \{(b_i, b_j) \mid (b_1, \dots, b_n) \in R\}$. A relation R is called 2-decomposable if, for any tuple $\mathbf{b} = (b_1, \dots, b_n) \in B^n$, we have $\mathbf{b} \in R$ if and only if $(b_i, b_j) \in \text{pr}_{i,j} R$ for all $1 \leq i < j \leq n$. It is well known (see, e.g., [15]) that any relation invariant under a majority operation is 2-decomposable.

Theorem 1. *Every majority operation on a k -element set has $(3k+2)$ -pathwidth duality.*

Proof idea: The following simple observation will be often used in the proof. Assume that we want to show that $\mathbf{A} \rightarrow \mathbf{B}$ for some fixed structures \mathbf{A} and \mathbf{B} . Then, if we have a structure \mathbf{C} such that $h : \mathbf{C}, c_1, \dots, c_n \rightarrow \mathbf{A}, a_1, \dots, a_n$ for some $a_1, \dots, a_n \in A$, then the possible values b_1, \dots, b_n taken by a_1, \dots, a_n , respectively, under a homomorphism from \mathbf{A} to \mathbf{B} must satisfy $\mathbf{C}, c_1, \dots, c_n \rightarrow \mathbf{B}, b_1, \dots, b_n$. This is because a composition of h with any homomorphism from \mathbf{A} to \mathbf{B} would be a homomorphism from \mathbf{C} to \mathbf{B} with this property.

We now explain the general strategy of our proof. First, we use 2-decomposability to reduce the situation to the case when all relations in structures under consideration are at most binary. Next, we prove that all structures of certain pathwidth that do not homomorphically map to \mathbf{B} form an obstruction set \mathcal{O} for \mathbf{B} . For this, we fix an arbitrary structure \mathbf{A} not admitting a homomorphism from any structure from \mathcal{O} , and use the above observation to reduce, for any $a \in A$, the set of values in B to which a can possibly be mapped by a homomorphism from \mathbf{A} to \mathbf{B} . Then, we choose values for elements in A (one by one in a certain order) from those reduced sets, while further reducing these sets at each step. Finally, we show that the obtained mapping $A \rightarrow B$ is homomorphism from \mathbf{A} to \mathbf{B} . The majority polymorphism is used to guarantee that this “greedy” approach never makes any reduced set empty. ■

Proof. (of Theorem 1). Fix a set B with $|B| = k$, and fix a majority operation ϕ on B . Call a structure *binary* if it has at most binary relations. The proof will use several lemmas.

Lemma 2. *The operation ϕ has $(3k+2)$ -pathwidth duality if every binary structure with polymorphism ϕ has $(3k+1, 3k+2)$ -pathwidth duality.*

Proof. Consider a new vocabulary τ' obtained from τ as follows: for every relation symbol R of arity $n \geq 3$, replace R by $\frac{n(n-1)}{2}$ binary relation symbols $R_{i,j}$, $1 \leq i < j \leq n$. Let \mathbf{C} be an arbitrary τ -structure. We transform it to an τ' -structure \mathbf{C}_{bin} as follows. Every at most binary relation in \mathbf{C} remains

unchanged in \mathbf{C}_{bin} . Every n -ary, $n \geq 3$, relation $R^{\mathbf{C}}$ is replaced, naturally, by relations $\text{pr}_{i,j} R^{\mathbf{C}}$, $1 \leq i < j \leq n$.

Since every relation in \mathbf{B} is 2-decomposable, it is easy to check that, for every τ -structure \mathbf{A} , $\mathbf{A} \rightarrow \mathbf{B}$ if and only if $\mathbf{A}_{bin} \rightarrow \mathbf{B}_{bin}$. It is well known, and also easy to check, that ϕ is also a polymorphism of the structure \mathbf{B}_{bin} . We need to show that \mathbf{B} has $(3k+2)$ -pathwidth duality if \mathbf{B}_{bin} has $(3k+1, 3k+2)$ -pathwidth duality.

Assume that \mathbf{B}_{bin} has an obstruction set of pathwidth at most $(3k+1, 3k+2)$ and let $r \geq 3$ be the maximum arity among the relations of \mathbf{B} . Let \mathbf{A} be any structure such that $\mathbf{A} \rightarrow \mathbf{B}$. We shall prove that there exists a structure of \mathbf{D} of pathwidth $(3k+2, 3k+r)$ that homomorphically maps to \mathbf{A} but not to \mathbf{B} . This will show that all structures of pathwidth at most $(3k+2, 3k+r)$ that do not have a homomorphism to \mathbf{B} form an obstruction set for \mathbf{B} , i.e., \mathbf{B} has $(3k+2)$ -duality.

Since $\mathbf{A} \rightarrow \mathbf{B}$, we have $\mathbf{A}_{bin} \rightarrow \mathbf{B}_{bin}$. Consequently, there exists a τ' -structure \mathbf{C} of pathwidth at most $(3k+1, 3k+2)$ such that $\mathbf{C} \rightarrow \mathbf{A}_{bin}$, but $\mathbf{C} \not\rightarrow \mathbf{B}_{bin}$. We obtain \mathbf{D} from \mathbf{C} in the following way. For any at most binary relation symbol $R \in \tau$, let $R^{\mathbf{D}} = R^{\mathbf{C}}$. For every relation symbol $R_{i,j}$ and every pair $(u, v) \in R_{i,j}^{\mathbf{C}}$, we introduce a tuple (d_1, \dots, d_n) in $R^{\mathbf{D}}$ such that $d_i = u$, $d_j = v$ and the remaining elements in the tuple are new elements that are particular to this tuple. It is straightforward to check that $\mathbf{D} \rightarrow \mathbf{A}$ and $\mathbf{D} \not\rightarrow \mathbf{B}$. It only remains to show that \mathbf{D} has pathwidth at most $(3k+2, 3k+r)$. Let S_1, \dots, S_m be a path-decomposition of \mathbf{C} , with $|S_i| \leq 3k+2$ for all $1 \leq i \leq m$. We shall obtain from it a path-decomposition of \mathbf{D} . For each relation symbol $R_{i,j}$ and for each pair (u, v) in $R_{i,j}^{\mathbf{C}}$ we select one set S_l , containing $\{u, v\}$. If necessary, we extend the sequence S_1, \dots, S_m by making copies of sets S_l , in order to ensure that the same set (as a member of the sequence) is never selected twice. Finally, we construct a path-decomposition from the (extended) sequence S_1, \dots, S_m in the following way: If a set S_l has been selected and associated to a given tuple (u, v) of the relation, say $R_{i,j}$, we enlarge S_l so that it also contains all the new elements from $\{d_1, \dots, d_n\}$ that have been introduced when processing tuple (u, v) in the construction of \mathbf{D} . Notice that the size of S_l increases at most by $n-2 \leq r-2$. Moreover, since the elements added while processing different tuples were different, the intersections of neighbours in the obtained sequence have size at most $3k+2$ (which is the maximum size of a set in the path-decomposition for \mathbf{C}). It is fairly easy to check that the obtained sequence of sets obtained is indeed a path-decomposition of \mathbf{D} , every set in it has at most $3k+r$ elements, and the size of the intersection of neighbour sets is at most $3k+2$. The lemma is proved. \blacksquare

In the rest of the proof, we assume that τ contains only at most binary relation symbols. We will show that that class of structures \mathbf{C} of pathwidth at most $(3k+1, 3k+2)$ such that $\mathbf{C} \rightarrow \mathbf{B}$ is an obstruction set for \mathbf{B} . Let \mathbf{A} be an arbitrary τ -structure. If there is a structure \mathbf{C} of pathwidth at most $(3k+1, 3k+2)$ such that $\mathbf{C} \rightarrow \mathbf{A}$ and $\mathbf{C} \not\rightarrow \mathbf{B}$ then, clearly, there is no homomorphism from \mathbf{A} to \mathbf{B} . Assume now that every structure \mathbf{C} of pathwidth at most $(3k+1, 3k+2)$

that homomorphically maps to \mathbf{A} also homomorphically maps to \mathbf{B} , and show that $\mathbf{A} \rightarrow \mathbf{B}$.

Let \mathcal{C}_n be the class of all τ -structures of pathwidth at most $(3n - 1, 3n)$.

Let \mathbf{C} be a τ -structure, and c a fixed element of \mathbf{C} . By Lemma 1, the unary relation $A(\mathbf{C}, c) = \{b \in B \mid \mathbf{C}, c \rightarrow \mathbf{B}, b\}$ is invariant under ϕ . For an element $a \in A$ and a number $1 \leq q \leq k$, let $A_a^q = \bigcap A(\mathbf{C}, c)$ where the intersection is taken over all pairs (\mathbf{C}, c) such that $\mathbf{C} \in \mathcal{C}_q$, $c \in C$, and $\mathbf{C}, c \rightarrow \mathbf{A}, a$. Since the intersection of subalgebras is always a subalgebra, A_a^q is a subalgebra for all a and q . Later, we will show that it is always non-empty.

A *path* P on a given structure \mathbf{C} is any sequence c_1, \dots, c_t of (possibly repeated) elements of the universe of C . The path P is a *cycle* if $c_1 = c_t$.

Let $P = a_1, \dots, a_t$, and $Q = b_1, \dots, b_t$ be paths on \mathbf{A} and \mathbf{B} , respectively, of the same length. We will denote the mapping from $\{a, a'\}$ to $\{b, b'\}$ taking a to b and a' to b' by $a, a' \rightarrow b, b'$. If $t > 1$, we say that Q *supports* P if, for all $1 \leq i < t$, the mapping $a_i, a_{i+1} \rightarrow b_i, b_{i+1}$ is a partial homomorphism from \mathbf{A} to \mathbf{B} . For $t = 1$, we say that Q *supports* P if the mapping $a_1 \rightarrow b_1$ is a partial homomorphism from \mathbf{A} to \mathbf{B} . Observe that several occurrences of the same element in the path on \mathbf{A} need not be mapped to the same value in \mathbf{B} . If, for some $n \geq 0$, we have $b_j \in A_{a_j}^n$ for all $1 \leq j \leq t$ then we say that b_1, \dots, b_t *n-supports* P .

Lemma 3. *Let $1 \leq n \leq k$ and let $P = a_1, \dots, a_t$ be any path on \mathbf{A} . There exists a structure \mathbf{C} of pathwidth at most $(3n + 1, 3n + 2)$ and some elements c_1, \dots, c_t in \mathbf{C} such that $\mathbf{C}, c_1, \dots, c_t \rightarrow \mathbf{A}, a_1, \dots, a_t$, and, furthermore, every path b_1, \dots, b_t in \mathbf{B} such that $\mathbf{C}, c_1, \dots, c_t \rightarrow \mathbf{B}, b_1, \dots, b_t$ *n-supports* P .*

Proof. First, we construct a structure \mathbf{C}' in the following way. Initially, \mathbf{C}' has one element that we call c'_1 . For every $1 \leq i < t$, we do the following. If $a_i \neq a_{i+1}$ then we include in the universe of \mathbf{C}' a new element that we denote by c'_{i+1} . If $a_i = a_{i+1}$ then we simply put $c'_{i+1} = c'_i$. Then we add to \mathbf{C}' all necessary tuples to ensure that, for all $1 \leq i < t$, the mapping $a_i, a_{i+1} \rightarrow c'_i, c'_{i+1}$ is an isomorphism from $\mathbf{A}_{\{a_i, a_{i+1}\}}$ to $\mathbf{C}'_{\{c'_i, c'_{i+1}\}}$ (i.e., this mapping and its inverse are both homomorphisms). By the definition of \mathbf{C}' , we have that $\mathbf{C}', c'_1, \dots, c'_t \rightarrow \mathbf{A}, a_1, \dots, a_t$. Furthermore, observe that every path b_1, \dots, b_t such that $\mathbf{C}', c'_1, \dots, c'_t \rightarrow \mathbf{B}, b_1, \dots, b_t$ supports a_1, \dots, a_t . Indeed, for every $1 \leq i < t$, the mapping $a_i, a_{i+1} \rightarrow b_i, b_{i+1}$ can be obtained by composing the mappings $a_i, a_{i+1} \rightarrow c'_i, c'_{i+1}$, and $c'_i, c'_{i+1} \rightarrow b_i, b_{i+1}$. By construction of \mathbf{C}' , the first mapping is a partial homomorphism from \mathbf{A} to \mathbf{C}' . By assumption, the second mapping is a partial homomorphism from \mathbf{C}' to \mathbf{B} . This proves that $a_i, a_{i+1} \rightarrow b_i, b_{i+1}$ is a partial homomorphism from \mathbf{A} to \mathbf{B} , since the composition of partial homomorphisms is again a partial homomorphism. It is easy to see that S_1, \dots, S_{t-1} , with $S_i = \{c'_i, c'_{i+1}\}$ is a path-decomposition of \mathbf{C}' of width $(1, 2)$.

For every $1 \leq i \leq t$ and for every $b \notin A_{a_i}^n$, there exists some structure $\mathbf{D}^{i,b}$ in \mathcal{C}_n and some element $d^{i,b}$ in \mathbf{D} such that $\mathbf{D}^{i,b}, d^{i,b} \rightarrow \mathbf{A}, a_i$ and $\mathbf{D}^{i,b}, d^{i,b} \rightarrow \mathbf{B}, b$. The structure \mathbf{C} is obtained by taking the disjoint union of the structures $\mathbf{D}^{i,b}$, for all $1 \leq i \leq t$ and $b \notin A_{a_i}^n$, and of \mathbf{C}' , and then identifying every $d^{i,b}$ with

c'_i . Finally, set $c_i = c'_i$ for all $1 \leq i \leq t$. Let us verify that $\mathbf{C}, c_1, \dots, c_t$ have the required properties. First, we shall prove that $\mathbf{C}, c_1, \dots, c_t \rightarrow a_1, \dots, a_t$. Let h be the partial mapping, with domain $\{c_1, \dots, c_t\}$, that sends c_i to a_i for all i . We know that $\mathbf{C}', c'_1, \dots, c'_t \rightarrow \mathbf{A}, a_1, \dots, a_t$ and consequently, h preserves all relations in \mathbf{C} restricted to $\{c_1, \dots, c_t\}$. (Note that $\mathbf{C}|_{\{c_1, \dots, c_t\}}$ may have a unary relation $\{c_i\}$ which was not in \mathbf{C}' , but was $\{d^{i,b}\}$ before identifying $d^{i,b}$ with c'_i – this relation $\{c_i\}$ is preserved because $\mathbf{D}^{i,b}, d^{i,b} \rightarrow \mathbf{A}, a_i$). Consider now any structure $\mathbf{D}^{i,b}$ attached to \mathbf{C}' when forming \mathbf{C} . Since $\mathbf{D}^{i,b}, d^{i,b} \rightarrow \mathbf{A}, a_i$ and $d^{i,b}$ is identified with $c'_i = c_i$ we can extend the homomorphism h to the elements of $\mathbf{D}^{i,b}$ so that h also preserves all relations in $\mathbf{D}^{i,b}$. Hence, we have proved that $\mathbf{C}, c_1, \dots, c_t \rightarrow \mathbf{A}, a_1, \dots, a_t$.

Now, let b_1, \dots, b_t be any path such that $\mathbf{C}, c_1, \dots, c_t \rightarrow \mathbf{B}, b_1, \dots, b_t$. Consequently, $\mathbf{C}', c'_1, \dots, c'_t \rightarrow \mathbf{B}, b_1, \dots, b_t$. It has been shown earlier that b_1, \dots, b_t supports a_1, \dots, a_t . Now let us show that $b_i \in A_{a_i}^n$ for all $1 \leq i \leq t$. Indeed, if $b_i \notin A_{a_i}^n$ then the mapping $c_i \rightarrow b_i$ cannot be extended to a homomorphism $\mathbf{D}^{i,b}, d^{i,b} \rightarrow \mathbf{B}, b_i$. Since $\mathbf{D}^{i,b}, d^{i,b} \rightarrow \mathbf{C}, c_i$, it follows that $\mathbf{C}, c_i \not\rightarrow \mathbf{B}, b_i$, a contradiction. Consequently, $b_i \in A_{a_i}^n$. It only remains to show that \mathbf{C} has the right pathwidth. We have seen above that there exists a path decomposition S_1, \dots, S_m of \mathbf{C}' of width $(1, 2)$.

For each $\mathbf{D}^{i,b}$ that we use in the construction of \mathbf{C} we select a set S_l containing c_i . If necessary, we extend the sequence S_1, \dots, S_m by making copies of sets S_l , in order to ensure that the same set (as a member of the sequence) is never selected twice. Finally, we construct a path decomposition of \mathbf{C} from the (extended) sequence S_1, \dots, S_m in the following way: If a set S_l has not been associated to any $\mathbf{D}^{i,b}$ then we leave as it is. Otherwise, we take a path-decomposition S'_1, \dots, S'_s of the structure $\mathbf{D}^{i,b}$ to which S_l has been associated and we replace S_l by the sequence $S_l \cup S'_1, \dots, S_l \cup S'_s$. It is fairly easy to verify that we obtain a path-decomposition of \mathbf{C} . Since each set in the path decomposition of each $\mathbf{D}^{i,b}$ can be assumed to have cardinality at most $3n$, the width of the path-decomposition of \mathbf{C} is at most $(3n + 1, 3n + 2)$. ■

From Lemma 3, we obtain the following corollaries.

Corollary 1. *Every path a_1, \dots, a_t in \mathbf{A} is k -supported by some path in \mathbf{B} .*

Proof. Use Lemma 3 with a_1, \dots, a_t and $n = k$. Let \mathbf{C} and c_1, \dots, c_t be the structure and elements provided by Lemma 3. Since \mathbf{C} has pathwidth at most $(3k + 1, 3k + 2)$ and $\mathbf{C} \rightarrow \mathbf{A}$, then there exists a homomorphism h from \mathbf{C} to \mathbf{B} . Let $b_i = h(c_i)$, $1 \leq i \leq t$. Lemma 3 guarantees that the path b_1, \dots, b_t k -supports a_1, \dots, a_k . ■

Corollary 2. *$A_a^q \neq \emptyset$ for all $a \in A$ and $1 \leq q \leq k$.*

Proof. In order to prove the case $q = k$, simply use Corollary 1 with the one-element path $P = a$. The other cases follow because $A_a^k \subseteq A_a^q$. ■

Corollary 3. *Let $q \geq 1$. Let a be any element of \mathbf{A} and let b be any element in A_a^q . Every cycle $a = a_1, \dots, a_t = a$ in \mathbf{A} is $(q-1)$ -supported by a cycle b_1, \dots, b_t in \mathbf{B} , with $b_1 = b_t = b$.*

Proof. Appeal to Lemma 3 with a_1, \dots, a_t and $n = q - 1$. Let \mathbf{C} and c_1, \dots, c_t be the structure and elements provided by Lemma 3. Let \mathbf{C}' be the structure obtained from \mathbf{C} by identifying the elements c_1 and c_t into a single element c . It is easy to see that identifying two elements can increase the pathwidth by at most one and hence, \mathbf{C}' belongs to \mathcal{C}_q . Since $a_1 = a_t$ we have that $\mathbf{C}', c \rightarrow \mathbf{A}, a$. By the definition of A_a^q , we have that $\mathbf{C}', c \rightarrow \mathbf{B}, b$. In consequence, there exists some b_1, \dots, b_m with $b_1 = b_m = b$ such that $\mathbf{C}, c_1, \dots, c_t \rightarrow b_1, \dots, b_t$ with $b = b_1 = b_t$. The properties of \mathbf{C} guarantee that b_1, \dots, b_t $(q-1)$ -supports a_1, \dots, a_t . ■

We will refer to Corollary 3 as to the *cycle q -condition*.

Note that, for every n , \mathcal{C}_n is a subclass of \mathcal{C}_{n+1} . This implies that, for all $a \in A$, we have $A_a^k \subseteq A_a^{k-1} \subseteq \dots \subseteq A_a^1 \subseteq A_a^0 = B$. Therefore, one of these inclusions is an equality, or else A_a^k is a singleton. It follows that, for every a , there is a smallest q such that A_a^q is an ideal in A_a^{q-1} . Call this number q_a . Assume that $|A| = m$ and order the elements of \mathbf{A} so that $q_{a_1} \geq q_{a_2} \geq \dots \geq q_{a_m}$. For simplicity, we will denote q_{a_i} by q_i . For $1 \leq i \leq m$, set $X_i = A_{a_i}^{q_i-1}$ and $Y_i = A_{a_i}^{q_i}$. Recall that Y_i is an ideal in X_i for all i .

Note that, for all $a_i \in A$, we have that $X_i \supseteq A_{a_i}^{k-1} \supseteq A_{a_i}^k$ and $Y_i \supseteq A_{a_i}^k$. By Corollary 1, we have the following property: every path a_{i_1}, \dots, a_{i_t} in \mathbf{A} is supported by a path b_{i_1}, \dots, b_{i_t} in \mathbf{B} such that $b_{i_1} \in Y_{i_1}, b_{i_t} \in Y_{i_t}$ and $b_{i_j} \in X_{i_j}$ for all $1 < j < t$. We call this property the *path 0-condition*.

We will now show how to choose elements b_1^*, \dots, b_m^* , in order, such that the mapping $h : A \rightarrow B$ with $h(a_i) = b_i^*$, $1 \leq i \leq m$, is a homomorphism from \mathbf{A} to \mathbf{B} . When we choose an element b_i^* , we will reset X_i and Y_i so that $X_i = B$ and $Y_i = \{b_i^*\}$. Note that we will always maintain the property that Y_i is an ideal in X_i , for all i .

Assume now that we have chosen b_1^*, \dots, b_r^* where $0 \leq r \leq m$. (The case $r = 0$ corresponds to the initial situation where no b_i^* is chosen yet). Then we have that $X_i = B$ and $Y_i = \{b_i^*\}$ for all $1 \leq i \leq r$. If, in the definition of the path 0-condition, we replace old X_i and Y_i , $1 \leq i \leq r$, with the new ones (i.e., $X_i = B$ and $Y_i = \{b_i^*\}$), then we will call the resulting condition the *path r -condition*.

We have shown above that the path 0-condition holds. Assume now that $l \geq 1$, and we have b_1^*, \dots, b_{l-1}^* such that the path $(l-1)$ -condition holds. Our goal now is to show that one can choose b_l^* so that the path l -condition holds. This will allow us to continue this process, and in the end, to prove that the obtained values b_i^* will indeed produce a homomorphism from \mathbf{A} to \mathbf{B} .

We will need some lemmas in order to choose b_l^* . Let $P = a_{i_1}, \dots, a_{i_t}$ be any path on \mathbf{A} . Consider a binary relation R_P consisting of all pairs (b, b') such that there exists a path b_{i_1}, \dots, b_{i_t} that supports P , with $b_{i_j} \in X_{i_j}$ for all $1 \leq j \leq t$, and also $b = b_{i_1}$ and $b' = b_{i_t}$.

Lemma 4. *The relation R_P is invariant under ϕ .*

Proof. We will prove the lemma by induction on t . It is well known, and easy to see, that the direct product of two subalgebras is invariant under ϕ , and intersections and compositions of binary relations invariant under ϕ are also invariant.

If $t = 1$ then $R_P = \{(b, b) \mid \mathbf{A} \upharpoonright_{\{a_l\}}, a_l \rightarrow \mathbf{B}, b\} \cap \{(b, b) \mid b \in X_l\}$. Since both X_l and $\{b \in B \mid \mathbf{A}, a_l \rightarrow \mathbf{B}, b\}$ are subalgebras, it follows that R_P is invariant. If $t = 2$ then $R_P = \{(b_{i_1}, b_{i_2}) \mid \mathbf{A} \upharpoonright_{\{a_{i_1}, a_{i_2}\}}, a_{i_1}, a_{i_2} \rightarrow \mathbf{B}, b_{i_1}, b_{i_2}\} \cap (X_{i_1} \times X_{i_2})$, so it is invariant. Let $t \geq 3$, and assume that the lemma holds for all shorter paths. Take paths $P' = a_{i_1}, \dots, a_{i_{t-1}}$ and $P'' = a_{i_{t-1}}, a_{i_t}$. It is easy to see that $R_P = R_{P'} \circ R_{P''}$ is the composition of $R_{P'}$ and $R_{P''}$. By inductive assumption, $R_{P'}$ and $R_{P''}$ are invariant. Then so is R_P . ■

Let $P = a_{i_1}, \dots, a_{i_t}$ be any path on \mathbf{A} where $a_{i_1} = a_l$. Let $U_P = \{b \in Y_l \mid (b, b') \in R_P \text{ for some } b' \in Y_{i_t}\}$. By the path $(l-1)$ -condition, U_P is non-empty.

Lemma 5. U_P is an ideal in Y_l .

Proof. We prove first that, for any $z \in Y_l$, there is $z' \in X_{i_t}$ such that $(z, z') \in R_P$. Consider the cycle $a_{i_1}, \dots, a_{i_t}, a_{i_1}$ obtained by adding a_{i_1} at the end of P .

Since $z \in Y_l$, then, by the cycle q_l -condition, the cycle $a_{i_1}, \dots, a_{i_t}, a_{i_1}$ is supported by a cycle $b_{i_1}, \dots, b_{i_t}, b_{i_1}$ with $b_{i_j} \in A_{a_{i_j}}^{q_l-1}$ for all $1 < j \leq t$ and $b_{i_1} = z$. Note that, due to the ordering of the elements of \mathbf{A} , and to the fact that $X_i = B$ for $1 \leq i \leq l-1$, we have that $A_i^{q_l-1} \subseteq X_i$ for all $1 \leq i \leq m$. Now let $z' = b_{i_t} (\in X_{i_t})$. By the definition of R_P , we get $(z, z') \in R_P$.

Now fix any x, y, z in Y_l such that (at least) two of these elements are in U_P , say $x, y \in U_P$. We show that $\phi(x, y, z) \in U_P$. We have shown above that we have $(z, z') \in R_P$ for some $z' \in X_{i_t}$. By the definition of U_P , there exists some tuple $(x, x') \in R_P$ with $x' \in Y_{i_t}$. Similarly, there exist some tuple $(y, y') \in R_P$ with $y' \in Y_{i_t}$. Since R_P is invariant under ϕ , the tuple $(\phi(x, y, z), \phi(x', y', z'))$ belongs to R_P . Since $x', y' \in Y_{i_t}$, $z' \in X_{i_t}$ and Y_{i_t} is an ideal in X_{i_t} , we conclude that $\phi(x', y', z') \in Y_{i_t}$. Consequently, $\phi(x, y, z) \in U_P$. ■

Lemma 6. If I_1, \dots, I_n , ($n \geq 2$), are ideals in a subalgebra U such that $I_i \cap I_j \neq \emptyset$ for all i, j , then we have $\bigcap_{1 \leq i \leq n} I_i \neq \emptyset$.

Proof. We prove the claim by induction on n . The base case $n = 2$ is trivial. Assume the claim holds for $n-1$ ideals and prove it for n . For $i = 1, 2, 3$, let J_i be the intersection of all ideals I_1, \dots, I_n except I_i . By inductive assumption, all three ideals J_i are non-empty. Choose $x \in J_1$, $y \in J_2$, and $z \in J_3$. For any ideal I_i , $1 \leq i \leq n$, at least two of x, y, z belong to I_i . It follows that $\phi(x, y, z) \in \bigcap_{1 \leq i \leq n} I_i$, and so $\bigcap_{1 \leq i \leq n} I_i \neq \emptyset$. ■

Let U_l be the intersection of all ideals of the form U_P , over all paths P starting at a_l .

Lemma 7. The set U_l is a non-empty subset of Y_l .

Proof. Since \mathbf{B} is finite, it is sufficient to prove non-emptiness for the intersection of any finite number of ideals of the form U_P . Furthermore, by Lemma 6, it is sufficient to show that the intersection is non-empty for any pair of such ideals.

Let $P = a_{i_1}, \dots, a_{i_t}$ and $Q = a_{g_1}, \dots, a_{g_s}$ be two arbitrary paths on \mathbf{A} with $a_{i_1} = a_{g_1} = a_l$. We need to show that $U_P \cap U_Q \neq \emptyset$.

Consider the following ternary relation R' on B : a triple (b, b', b'') belongs to R' if and only if both $(b, b') \in R_P$ and $(b, b'') \in R_Q$. Since both R_P and R_Q are invariant under ϕ , it is easy to verify that R' is invariant as well.

Consider the path $a_{i_1}, \dots, a_{i_2}, a_{i_1}, a_{g_2}, \dots, a_{g_s}$ on \mathbf{A} . By applying the path $(l-1)$ -condition to this path, we obtain that there exist $x_1 \in X_{i_1} (= X_l)$, $y_1 \in Y_{i_2}$ and $z_1 \in Y_{g_s}$ such that $(x_1, y_1, z_1) \in R'$. By applying the path $(l-1)$ -condition to P , we can obtain elements $x_2 \in Y_{i_1} (= Y_l)$ and $y_2 \in Y_{i_2}$ such that $(x_2, y_2) \in U_P$. Furthermore, since $x_2 \in Y_{i_1} = Y_l = Y_{g_1}$, we can, as in the (first part of the) proof of Lemma 5, find $z_2 \in X_{g_s}$ such that $(x_2, z_2) \in U_Q$. Hence, we have $(x_2, y_2, z_2) \in R'$. By symmetry, there is a triple $(x_3, y_3, z_3) \in R'$ such that $x_3 \in Y_{i_1} (= Y_l)$, $y_3 \in X_{i_2}$, and $z_3 \in Y_{g_s}$. Notice that, in each coordinate, these triples have at least two elements from the corresponding ideal Y_{i_j} . Now let $u_x = \phi(x_1, x_2, x_3)$, $u_y = \phi(y_1, y_2, y_3)$, and $u_z = \phi(z_1, z_2, z_3)$. Since R' is invariant, we have $(u_x, u_y, u_z) \in R'$. Since at least two of the x_i 's belong to Y_{i_1} , we have that $u_x \in Y_{i_1}$. Similarly, we have $u_y \in Y_{i_2}$ and $u_z \in Y_{g_s}$. Thus, we have $(u_x, u_y) \in R_P$ and $(u_x, u_y) \in R_Q$, which implies that $u_x \in U_P \cap U_Q$. ■

Now let b_l^* be an arbitrary element from U_l , and set $X_l = B$ and $Y_l = \{b_l^*\}$.

Lemma 8. *With b_l^* chosen as above, the path l -condition holds.*

Proof. Take any arbitrary path a_{i_1}, \dots, a_{i_t} on \mathbf{A} . We need to show that the path is supported by some path b_{i_1}, \dots, b_{i_t} on \mathbf{B} with $b_{i_j} \in X_{i_j}$ for all $1 < j < t$ and $b_{i_1} \in Y_{i_1}, b_{i_t} \in Y_{i_t}$.

If none of the elements a_{i_1}, \dots, a_{i_t} is a_l then we have the required path on \mathbf{B} by the path $(l-1)$ -condition, since the only difference between the path $(l-1)$ - and l -conditions is the definition of X_l and Y_l . Suppose now that the sequence a_{i_1}, \dots, a_{i_t} contains at least one occurrence of a_l .

Let $a_{i_{\min}}$ and $a_{i_{\max}}$ the first and last occurrence, respectively, of a_l in P . Let us consider the subpath $Q = a_{i_{\max}}, \dots, a_{i_t}$. Since $b_l^* \in U_l$ we can infer that $b_l^* \in U_Q$. Consequently, there exists some path $b_{i_{\max}}, \dots, b_{i_t}$ in \mathbf{B} supporting $a_{i_{\max}}, \dots, a_{i_t}$ with $b_{i_{\max}} = b_l^*, b_{i_t} \in Y_{i_t}$ and $b_{i_j} \in X_{i_j}$ for all $i_{\max} < j < t$. Similarly, there exists some path $b_{i_1}, \dots, b_{i_{\min}}$ in \mathbf{B} supporting $a_{i_1}, \dots, a_{i_{\min}}$ with $b_{i_1} \in Y_{i_1}, b_{i_{\min}} = b_l^*$ and $b_{i_j} \in X_{i_j}$ for all $1 < j < i_{\min}$. Let us consider now the cycle $a_{i_{\min}}, \dots, a_{i_{\max}}$. If $i_{\min} < i_{\max}$ (i.e., more than one occurrence of a_l) then, by the q_l -cycle condition, $a_{i_{\min}}, \dots, a_{i_{\max}}$ is supported by a cycle $b_{i_{\min}}, \dots, b_{i_{\max}}$ in \mathbf{B} with $b_{i_{\min}} = b_{i_{\max}} = b_l^*$ and $b_{i_j} \in A_{i_j}^{q_l-1}$ for all $i_{\min} < j < i_{\max}$. Due to the ordering of the elements of \mathbf{A} and to the fact that $X_i = B$ for $1 \leq i \leq l$, we have that $A_i^{q_l-1} \subseteq X_i$ for all $1 \leq i \leq m$. Consequently, the path $b_{i_1}, \dots, b_{i_{\min}}, \dots, b_{i_{\max}}, \dots, b_{i_t}$ supports P and satisfies $b_{i_1} \in Y_{i_1}, b_{i_t} \in Y_{i_t}$ and $b_{i_j} \in X_{i_j}$ for all $1 < j < t$. The lemma is proved. ■

We have shown that the path 0-condition holds, and, assuming that we can choose b_1^*, \dots, b_{l-1}^* , $l \geq 1$ so that the path $(l-1)$ -condition holds, we have shown that it is possible to choose b_l^* so that the path l -condition holds for $b_1^*, \dots, b_{l-1}^*, b_l^*$. Hence, the path m -condition holds. (Recall that $A = \{a_1, \dots, a_m\}$).

Lemma 9. *The mapping $h : A \rightarrow B$ such that $h(a_i) = b_i^*$ for all $1 \leq i \leq m$ is a homomorphism from \mathbf{A} to \mathbf{B} .*

Proof. Recall that every relation in \mathbf{A} is at most binary. Let R be any binary relation symbol in τ and let (a_i, a_j) be any tuple in $R^{\mathbf{A}}$. Consider the path a_i, a_j on \mathbf{A} . The path m -condition guarantees that b_i^*, b_j^* supports it. Consequently, the mapping h restricted to a_i, a_j is a partial homomorphism. Hence, $(b_i^*, b_j^*) \in R^{\mathbf{B}}$. The proof for unary relations is similar. ■

Theorem 1 is proved. ■

Corollary 4. *If a structure \mathbf{B} has a majority polymorphism then $\text{CSP}(\mathbf{B})$ is in complexity class \mathbf{NL} .*

Proof. By Proposition 3 [7], $\text{CSP}(\mathbf{B})$ is in \mathbf{NL} for any structure \mathbf{B} with bounded pathwidth duality. Now the result follows from Theorem 1. ■

It was shown in [10] that any structure \mathbf{B} with a majority polymorphism has 2-treewidth duality, while our Theorem 1 states that the parameter j in the j -pathwidth duality for such structures grows linearly with the size of the base set of the structure. In the next two sections we show that this linear growth is in fact unavoidable.

4 Pebble-Relation Games

In this section, we describe pebble-relation games, introduced in [6, 7], which can be used to characterise bounded pathwidth duality. We will use these games in next section to prove the existence of binary structures which, for a given $n \geq 1$, have $(6n+2)$ -pathwidth duality, but not n -pathwidth duality.

Let S_1 and S_2 be two sets. We define a *relation* T with domain S_1 and range S_2 as a collection of functions with domain S_1 and range S_2 . Some confusion might arise from the fact that generally (and in this paper) the name relation is used with another meaning; for example, an r -ary relation over B is a subset of B^r . Both concepts are perfectly consistent, since an r -ary relation over B is, indeed, a relation in the new sense, with domain $\{1, \dots, r\}$ and range B . We will assume by convention that for every set B there exists one mapping $\lambda : \emptyset \rightarrow B$.

Let f be a function with domain S_1 and range S_2 , and let S'_1 be a subset of its domain S_1 . We will denote by $f|_{S'_1}$ the restriction of f to S'_1 . Similarly, let T be a relation with domain S_1 and range S_2 , and let S'_1 be a subset of its domain S_1 . We will denote by $T|_{S'_1}$ the relation with domain S'_1 and range S_2 .

that contains $f|_{S_1^c}$ for every $f \in T$. For every relation T we denote by $\text{dom}(T)$ the domain of T . We have two relations with domain \emptyset : the relation $\{\lambda\}$ and the relation \emptyset .

Let $0 \leq j \leq k$ be non-negative integers and let \mathbf{A} and \mathbf{B} be τ -structures. The (j, k) -pebble-relation ((j, k) -PR) game on \mathbf{A} and \mathbf{B} is played between two players, the *Spoiler* and the *Duplicator*. A configuration of the game consists of a relation T with domain $I = \{a_1, \dots, a_{k'}\} \subseteq A$, $k' \leq k$ and range B such that every function f in T is a homomorphism from $\mathbf{A}|_I$ to \mathbf{B} .

Initially $I = \emptyset$ and T contains the (unique) homomorphism from $\mathbf{A}|_{\emptyset}$ to \mathbf{B} , that is, λ . Each round of the game consists of a move from the Spoiler and a move from the Duplicator. Intuitively, the Spoiler has control on the domain I of T , which can be regarded as placing some pebbles on the elements of A that constitute I , whereas the Duplicator decides the content of T after the domain I has been set by the Spoiler. There are two types of rounds: *shrinking* rounds and *blowing* rounds.

Let T^n be the configuration after the n -th round. The spoiler decides whether the following round is a blowing or shrinking round.

- If the $(n+1)$ -th round is a shrinking round, the Spoiler sets I^{n+1} (the domain of T^{n+1}) to be a subset of the domain I^n of T^n . The Duplicator responds by restricting every function in T^n onto the subdomain defined by I^{n+1} , that is, $T^{n+1} = T^n|_{I^{n+1}}$.
- A blowing round only can be performed if $|I^n| \leq j$. In this case the Spoiler sets I^{n+1} to be a superset of I^n with $|I^{n+1}| \leq k$. The duplicator responds by providing a T^{n+1} with domain I^{n+1} such that $T^n|_{I^n} \subseteq T^{n+1}$. That is, T^{n+1} should contain some extensions of functions in T^n over the domain I^{n+1} (recall that any such extension must be a homomorphism from $\mathbf{A}|_{I^{n+1}}$ to \mathbf{B}).

The Spoiler wins the game if the response of the Duplicator sets T^{n+1} to \emptyset , i.e., the Duplicator could not extend successfully any of the functions. Otherwise, the game resumes. The Duplicator wins the game if he has an strategy that allows him to continue playing “forever”, i.e., if the Spoiler can never win a round of the game.

We denote by $\text{hom}(\mathbf{A}, \mathbf{B})$ the set of all homomorphisms from \mathbf{A} to \mathbf{B} .

Now, we will present an algebraic characterization of the (j, k) -PR game.

Definition 5. *Let $0 \leq j < k$ be non-negative integers and let \mathbf{A} and \mathbf{B} be τ -structures. We say that the Duplicator has a winning strategy for the (j, k) -pebble-relation game on \mathbf{A} and \mathbf{B} if there is a nonempty family \mathcal{H} of relations such that:*

- (a) every relation T has range B and domain I for some $I \subseteq A$ with $|I| \leq k$.
- (b) for every relation T in \mathcal{H} with domain I , $\emptyset \neq T$ and $T \subseteq \text{hom}(\mathbf{A}|_I, \mathbf{B})$
- (c) \mathcal{H} is closed under restrictions: for every T in \mathcal{H} with domain I and every $I' \subseteq I$, we have that $T|_{I'} \in \mathcal{H}$.

- (d) \mathcal{H} has the (j, k) -forth property: for every relation T in \mathcal{H} with domain I with $|I| \leq j$ and every superset I' of I with $|I'| \leq k$, there exists some relation T' in \mathcal{H} with domain I' such that $T'_I \subseteq T$.

The intuition behind the definition of a winning strategy is that every relation in a winning strategy corresponds to a winning configuration for the Duplicator in the game.

The following result will be most useful.

Theorem 2. [7] *Let $0 \leq j < k$ be non-negative integers and let \mathbf{A} and \mathbf{B} be finite τ -structures. Then the two following conditions are equivalent:*

- *Duplicator has a winning strategy for the (j, k) -PR game.*
- *For every structure \mathbf{C} of pathwidth at most (j, k) , if $\mathbf{C} \rightarrow \mathbf{A}$ then $\mathbf{C} \rightarrow \mathbf{B}$.*

Indeed, by virtue of Theorem 2, in order to prove that a certain structure \mathbf{B} does not have (j, k) -pathwidth duality, we only need to provide an structure $\mathbf{A} \notin \text{CSP}(\mathbf{B})$ and provide a winning strategy for the Duplicator in the (j, k) -PR game.

5 Pathwidth Hierarchy Does Not Collapse

Recall that we call a structure binary if all of its relations are at most binary.

Theorem 3. *For every $n \geq 1$, there exists a binary structure with $2n$ elements which has a majority polymorphism, but does not have n -pathwidth duality.*

Proof. Let $n \geq 1$, and let B_n be the set $\{1, \dots, n\} \times \{1, 2\}$. The *level* of an element (i, j) of B_n is defined to be its first coordinate i . For every $1 \leq k \leq n$, R_n^k is a binary symmetric relation on B_n that consists of all pairs $((i, j), (i', j'))$ satisfying at least one of the following conditions

- $i > k, i = i', j = j'$
- $i = i' = k, j \neq j'$,
- $i < k$ and $i' \leq k$,
- $i \leq k$ and $i' < k$.

Let f_n be the ternary majority operation on B_n that returns, when the majority rule does not apply (i.e., when there is no repetition among the arguments), the first (from left to right) argument from the lowest level. Let us prove that R_n^k is invariant under f_n for every k . Take two triples (a_1, b_1, c_1) and (a_2, b_2, c_2) of elements of B_n such that $(a_1, a_2), (b_1, b_2), (c_1, c_2) \in R_n^k$ and show that $(d_1, d_2) = (f_n(a_1, b_1, c_1), f_n(a_2, b_2, c_2)) \in R_n^k$. Let k_1 be the lowest level of an element in the first triple and k_2 in the second. Clearly, if the majority rule applies to both triples then (d_1, d_2) is one of $(a_1, a_2), (b_1, b_2), (c_1, c_2)$, and so belongs to R_n^k .

Assume that the majority rule applies to exactly one of the triples, say, the first one. The level of the repeated element in the first triple cannot be greater

than k because then the corresponding elements in the second triple would also have to coincide. If the level of the repeated element is at most $k - 1$ then we have $k_2 \leq k$, and so $(d_1, d_2) \in R_n^k$. If the level of the repeated element is exactly k then, since there is no repetition in the second triple, we have $k_2 < k$ and so $(d_1, d_2) \in R_n^k$.

If neither triple has a repetition, then the proof is very similar when at least one k_1, k_2 is not equal to k . If $k_1 = k_2 = k$ then the first argument of level k in the first triple is (k, i) , while the the first argument of level k in the second triple is (k, j) for $i \neq j$, and so we have $(d_1, d_2) \in R_n^k$ again.

Let $\tau_n = \{P_n^1 \dots, P_n^n\}$ be the vocabulary that contains a binary relation symbol P_n^k for every $1 \leq k \leq n$. Let \mathbf{B}_n be the τ -structure with universe B_n and such that every P_n^k , $1 \leq k \leq n$, is interpreted as R_n^k .

We shall show that \mathbf{B}_n does not have n -pathwidth duality. For every $m > n$, we shall construct a structure \mathbf{A}_n^m that certifies that \mathbf{B}_n does not have (n, m) -pathwidth duality.

We define \mathbf{A}_n^m by induction on n . Let M be any odd number strictly larger than m . The universe of \mathbf{A}_1^m contains M elements x_0, \dots, x_{M-1} . For every $0 \leq i \leq M - 1$, we include the tuple x_i, x_{i+1} (here the addition is modulo M) in $(P_1^1)^{\mathbf{A}_1^m}$.

In order to construct \mathbf{A}_n^m , we first consider 3 copies $\mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3$ of \mathbf{A}_{n-1}^m (we can without loss of generality assume that they have disjoint sets of elements). Observe that every copy \mathbf{C}_i , $i \in \{1, 2, 3\}$ is a τ_{n-1} -structure. We transform it into a τ_n -structure \mathbf{C}'_i by merely setting $(P_n^k)^{\mathbf{C}'_i} = (P_{n-1}^k)^{\mathbf{C}_i}$ for all $1 \leq k < n$ and $(P_n^n)^{\mathbf{C}'_i} = \emptyset$. We call \mathbf{C}'_i a *modified copy*. Now we arbitrarily select one element in each of the modified copies. Let us denote them by y_1, y_2, y_3 . The structure \mathbf{A}_n^m is obtained by first computing the disjoint union of the three modified copies $\mathbf{C}'_1 \cup \mathbf{C}'_2 \cup \mathbf{C}'_3$, and then setting $(P_n^n)^{\mathbf{A}_n^m}$ to consist of the pairs $(y_1, y_2), (y_2, y_3), (y_3, y_1)$.

Note that the partial homomorphisms from \mathbf{C}_i to \mathbf{B}_{n-1} are precisely those partial homomorphisms from \mathbf{C}'_i to \mathbf{B}_n that do not have elements of level n (i.e., $(n, 1)$ and $(n, 2)$) in their images.

Let us first prove that \mathbf{A}_n^m is not homomorphic to \mathbf{B}_n , again by induction on n . If $n = 1$ then R_1^1 is the disequality relation (\neq) on the set $B_1 = \{(1, 1), (1, 2)\}$. Hence \mathbf{A}_1^m is not homomorphic to \mathbf{B}_1 because the only relation in \mathbf{A}_1^m is a cycle of odd length M . We shall prove now that $\mathbf{A}_n^m \not\rightarrow \mathbf{B}_n$ if $\mathbf{A}_{n-1}^m \not\rightarrow \mathbf{B}_{n-1}$. Let \mathbf{C}_i , $i \in \{1, 2, 3\}$ be any copy of \mathbf{A}_{n-1}^m used in the construction \mathbf{A}_n^m and let \mathbf{C}'_i be its modified copy. The set of homomorphisms from \mathbf{C}'_i to \mathbf{B}_n is easy to describe. Since \mathbf{C}_i is not homomorphic to \mathbf{B}_{n-1} , every homomorphism from \mathbf{C}'_i to \mathbf{B}_n must necessarily map at least one element of \mathbf{C}'_i to one of the new values $(n, 1)$ or $(n, 2)$. Since every relation R_n^k with $1 \leq k \leq n - 1$ forces the values of level n to be identical and \mathbf{C}'_i is connected, any other element in \mathbf{C}'_i has to take the same value. Consequently, each \mathbf{C}'_i has only two homomorphisms to \mathbf{B}_n : one of them sends all elements to $(n, 1)$ and the other to $(n, 2)$. Finally let us take into consideration the distinguished elements y_1, y_2, y_3 in \mathbf{A}_n^m . A homomorphism from \mathbf{A}_n^m to \mathbf{B}_n can map these elements only to the values of level n in \mathbf{B}_n . However,

these elements constitute a cycle of length 3 in $(P_n^m)^{\mathbf{A}_n^m}$, while R_n^m restricted to the values of level n is the disequality relation. Hence, $\mathbf{A}_n^m \not\rightarrow \mathbf{B}_n$.

It remains to show the Duplicator has a winning strategy \mathcal{H}_n^m for the (n, m) -PR game on \mathbf{A}_n^m and \mathbf{B}_n . Again, we prove this by induction on n . Let us consider first the case $n = 1$. Recall that in this case \mathbf{A}_1^m is essentially a cycle with M vertices and that R_1^1 is the disequality relation. It is an easy task to find a winning strategy for the Duplicator. Specifically, the winning strategy \mathcal{H}_1^m contains, for each set I with size at most m , a relation that contains precisely all partial homomorphisms from \mathbf{A}_1^m to \mathbf{B}_1 with domain I . Note that, since $m < M$, $\mathbf{A}_1^m|_I$ is a family of disjoint paths, and every partial homomorphism from $\mathbf{A}_1^m|_I$ to \mathbf{B}_1 can be seen as a 2-coloring of these paths. It is straightforward to check that \mathcal{H}_1^m is indeed a winning strategy.

We now show how to construct a winning strategy for the Duplicator for the (n, m) -PR game on \mathbf{A}_n^m and \mathbf{B}_n . Let \mathcal{H}^i , $i \in \{1, 2, 3\}$ be winning strategies for $(n-1, m)$ -PR game on \mathbf{C}_i and \mathbf{B}_{n-1} for the 3 copies $\mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3$ of \mathbf{A}_{n-1}^m that are used in the construction of \mathbf{A}_n^m .

Let I be any subset of \mathbf{A}_n^m , $|I| \leq m$, and let I_1, I_2, I_3 be subsets of I such that each set I_i , $i \in \{1, 2, 3\}$ contains precisely those elements of I that belong to \mathbf{C}'_i . (Note that some of the sets I_i may be empty). Let j_1, j_2 be different values in $\{1, 2, 3\}$ and j_3 be the remaining value. Let R_{j_1} be any relation of \mathcal{H}^{j_1} with domain I_{j_1} and similarly let R_{j_2} be any relation of \mathcal{H}^{j_2} with domain I_{j_2} .

We define $R = R(I, j_1, j_2, R_{j_1}, R_{j_2})$ to be the relation that contains all those partial homomorphisms \mathbf{t} with domain I that satisfy the following conditions:

- $\mathbf{t}|_{I_{j_1}} \in R_{j_1} \cup \{(\mathbf{n}, \mathbf{1}), (\mathbf{n}, \mathbf{2})\}$,
- $\mathbf{t}|_{I_{j_2}} \in R_{j_2} \cup \{(\mathbf{n}, \mathbf{1}), (\mathbf{n}, \mathbf{2})\}$, and
- $\mathbf{t}|_{I_{j_3}} \in \{(\mathbf{n}, \mathbf{1}), (\mathbf{n}, \mathbf{2})\}$.

where $(\mathbf{n}, \mathbf{1})$ (respectively, $(\mathbf{n}, \mathbf{2})$) denotes the mapping, with the corresponding domain, that maps all elements to $(n, 1)$ (respectively, $(n, 2)$).

We define \mathcal{H}_n^m to be the set that contains $R(I, j_1, j_2, R_{j_1}, R_{j_2})$ for all valid choices of I, j_1, j_2, R_{j_1} , and R_{j_2} .

We shall prove that \mathcal{H}_n^m is indeed a winning strategy. It is fairly easy to verify that, for every choice of I, j_1, j_2, R_{j_1} , and R_{j_2} , the relation $R = R(I, j_1, j_2, R_{j_1}, R_{j_2})$ is nonempty: for example, any mapping \mathbf{t} with $\mathbf{t}|_{I_{j_1}} \in R_{j_1}$, $\mathbf{t}|_{I_{j_2}} = (\mathbf{n}, \mathbf{1})$, and $\mathbf{t}|_{I_{j_3}} = (\mathbf{n}, \mathbf{2})$ is a partial homomorphism. Moreover, one can show by slightly modifying the previous example that the restriction of R to each I_{j_i} , $i \in \{1, 2, 3\}$ contains $(\mathbf{n}, \mathbf{1})$ and $(\mathbf{n}, \mathbf{2})$ (with the corresponding domain). We shall use this fact later.

It follows directly from the definitions that \mathcal{H}_n^m is closed under restrictions.

It only remains to show that \mathcal{H}_n^m has the (n, m) -forth property. Let $R = R(I, j_1, j_2, R_{j_1}, R_{j_2})$ be any relation in \mathcal{H}_n^m with $|I| \leq n$ and let I' be any superset of I with $|I'| \leq m$. We shall show that there exists some relation R' in \mathcal{H}_n^m with domain I' such that $R'|_I$ is contained in R . Let us consider two cases.

First let us assume that the cardinality of I_{j_1} or the cardinality of I_{j_2} is n . Assume without loss of generality that the former holds. Notice that in this case

$I = I_{j_1}$. We just set $R' = R(I', j_2, j_3, R'_{j_2}, R'_{j_3})$ where R'_{j_2} is any relation in \mathcal{H}^{j_2} with domain $I' \cap C'_{j_2}$ and R'_{j_3} is any relation in \mathcal{H}^{j_3} with domain $I' \cap C'_{j_3}$. The restriction of R' to I , which consists of $(\mathbf{n}, \mathbf{1})$ and $(\mathbf{n}, \mathbf{2})$ (with domain I), is contained in R .

Assume now that $|I_{j_1}| \leq n - 1$ and $|I_{j_2}| \leq n - 1$. By the $(n - 1, m)$ -forth property of \mathcal{H}^{j_1} , there exists a relation R'_{j_1} in \mathcal{H}^{j_1} with domain $I'_{j_1} (= I' \cap C'_{j_1})$ such that $R'_{j_1}|_{I_{j_1}} \subseteq R_{j_1}$. Similarly, by the $(n - 1, m)$ -forth property of \mathcal{H}^{j_2} , there exists a relation R'_{j_2} in \mathcal{H}^{j_2} with domain $I'_{j_2} (= I' \cap C'_{j_2})$ such that $R'_{j_2}|_{I_{j_2}} \subseteq R_{j_2}$. We obtain the required relation R' by setting $R' = R(I', j_1, j_2, R'_{j_1}, R'_{j_2})$. ■

Theorem 3 is of interest in the context of the so-called duality (or Datalog) hierarchies. It is an important open question whether, for any $n \geq 3$, there is a structure which has bounded treewidth duality, but not n -treewidth duality. For pathwidth dualities, this question can be answered in positive by using structures with only two elements (see Section 7.2 of [7]), but the arity of the relations in such structures would grow with n . Theorem 3 shows that, even when the arity of relations in structures is bounded by 2, there exist structures for which the parameter j of j -pathwidth duality is (necessarily) arbitrarily large.

6 Conclusion

We have shown that every relational structure with a majority polymorphism has bounded pathwidth duality, thus solving an open problem posed in [6, 7]. There are two natural extensions of this class of structures, for which it seems reasonable to (try to) prove the existence of bounded pathwidth duality.

One is the class of structures having a *near-unanimity* polymorphism, which is an n -ary ($n \geq 3$) operation f satisfying, for all x, y , the identities

$$f(y, x, x, \dots, x, x) = f(x, y, x, \dots, x, x) = \dots = f(x, x, x, \dots, x, y) = x.$$

Clearly, a majority operation is simply a ternary near-unanimity operation. It is known [10] that any structure with a near-unanimity polymorphism of arity $n + 1$ has n -treewidth duality. In [10], such structures were shown to have a special property called “bounded strict width”. The problem of whether such structures have bounded pathwidth duality was also mentioned in [6, 7].

The other class, which is known to properly extend the previous one, consists of structures that admit a sequence of *Jónsson* operations (as polymorphisms). For $k \geq 2$, a sequence of ternary operations p_i ($0 \leq i \leq k$) is called a sequence of *Jónsson* operations if the operations satisfy the following identities:

$$\begin{aligned} p_0(x, y, z) &= x \\ p_k(x, y, z) &= z \\ p_i(x, y, x) &= x \text{ for all } i \\ p_i(x, x, y) &= p_{i+1}(x, x, y) \text{ for all even } i \\ p_i(y, x, x) &= p_{i+1}(y, x, x) \text{ for all odd } i \end{aligned}$$

Such sequences of operations are studied in universal algebra in connection with the property of congruence-distributivity. Note that if such a sequence has three operations (i.e., $k = 2$) then p_1 is simply a majority operation. It is known [17] that a sequence of four Jónsson operations (as polymorphisms) guarantees bounded treewidth duality, but even this question is still open for such sequences with more than four operations.

Acknowledgements

This work was partially supported by the UK EPSRC grant EP/C54384X/1. Part of this work was done when both authors visited the Isaac Newton Institute for Mathematical Sciences, Cambridge. Financial support from the Institute is gratefully acknowledged. The first author is also supported by the MEC under the program "Ramon y Cajal", grants TIC 2002-04470-C03, TIC 2002-04019-C03, TIN 2004-04343, the EU PASCAL Network of Excellence IST-2002-506778, and the MODNET Marie Curie Research Training Network MRTN-CT-2004-512234. The second author is also supported by the UK EPSRC grant EP/C543831/1.

References

1. A. Atserias. On digraph coloring problems and treewidth duality. In *LICS'05*, pages 106–115, 2005.
2. A. Bulatov. Tractable conservative constraint satisfaction problems. In *LICS'03*, pages 321–330, 2003.
3. A. Bulatov. A graph of a relational structure and constraint satisfaction problems. In *LICS'04*, pages 448–457, 2004.
4. A. Bulatov, P. Jeavons, and A. Krokhin. Classifying complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.
5. 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.
6. V. Dalmau. Constraint satisfaction problems in non-deterministic logarithmic space. In *ICALP'02*, volume 2380 of *LNCS*, pages 414–425, 2002.
7. V. Dalmau. Linear Datalog and bounded path duality for relational structures. *Logical Methods in Computer Science*, 1, 2005. (electronic).
8. V. Dalmau, A. Krokhin, and B. Larose. First-order definable retraction problems for reflexive graphs and posets. In *LICS'04*, pages 232–241, 2004.
9. T. Feder, P. Hell, and J. Huang. Bi-arc graphs and the complexity of list homomorphisms. *Journal of Graph Theory*, 42:61–80, 2003.
10. T. Feder and M.Y. 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.
11. M. Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. In *FOCS'03*, pages 552–561, 2003.
12. P. Hell and J. Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.

13. P. Hell, J. Nešetřil, and X. Zhu. Duality and polynomial testing of tree homomorphisms. *Trans. Amer. Math. Soc.*, 348:147–156, 1996.
14. P. Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200:185–204, 1998.
15. P.G. Jeavons, D.A. Cohen, and M.C. Cooper. Constraints, consistency and closure. *Artificial Intelligence*, 101(1-2):251–265, 1998.
16. P.G. Jeavons, D.A. Cohen, and M. Gyssens. Closure properties of constraints. *Journal of the ACM*, 44:527–548, 1997.
17. E. Kiss and M. Valeriote. On tractability and congruence distributivity. In *LICS'06*, 2006. to appear.
18. Ph.G. Kolaitis and M.Y. Vardi. Conjunctive-query containment and constraint satisfaction. *Journal of Computer and System Sciences*, 61:302–332, 2000.
19. A. Krokhin, A. Bulatov, and P. Jeavons. The complexity of constraint satisfaction: an algebraic approach. In *Structural Theory of Automata, Semigroups, and Universal Algebra*, volume 207 of *NATO Science Series II: Math., Phys., Chem.*, pages 181–213. Springer Verlag, 2005.
20. A. Krokhin and B. Larose. Solving order constraints in logarithmic space. In *STACS'03*, volume 2607 of *LNCS*, pages 379–390, 2003.
21. B. Larose, C. Loten, and C. Tardif. A characterisation of first-order constraint satisfaction problems. In *LICS'06*, 2006. to appear.
22. B. Larose and L. Zádori. Bounded width problems and algebras. *Algebra Universalis*, 2006. to appear.
23. 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.