

Bounded width problems and algebras

BENOIT LAROSE AND LÁSZLÓ ZÁDORI

ABSTRACT. Let \mathcal{A} be finite relational structure of finite type, and let $CSP(\mathcal{A})$ denote the following decision problem: if \mathcal{I} is a given structure of the same type as \mathcal{A} , is there a homomorphism from \mathcal{I} to \mathcal{A} ? To each relational structure \mathcal{A} is associated naturally an algebra \mathbb{A} whose structure determines the complexity of the associated decision problem. We investigate those finite algebras arising from CSP's of so-called *bounded width*, i.e., for which local consistency algorithms effectively decide the problem. We show that if a CSP has bounded width then the variety generated by the associated algebra omits the Hobby–McKenzie types 1 and 2. This provides a method to prove that certain CSP's do not have bounded width. We give several applications, answering a question of Nešetřil and Zhu [26], by showing that various graph homomorphism problems do not have bounded width. Feder and Vardi [17] have shown that every CSP is polynomial-time equivalent to the retraction problem for a poset we call the *Feder–Vardi poset* of the structure. We show that, in the case where the structure has a single relation, if the retraction problem for the Feder–Vardi poset has bounded width then the CSP for the structure also has bounded width. This is used to exhibit a finite order-primal algebra whose variety admits type 2 but omits type 1 (provided $\mathbf{P} \neq \mathbf{NP}$).

1. Introduction

Throughout this paper the relational structures considered are finite and of finite type, i.e., we consider structures of the form $\mathcal{A} = \langle A; \theta_1, \dots, \theta_s \rangle$ where A is a finite non-empty set and the θ_i are finitary relations on A . As usual, a structure $\mathcal{B} = \langle B; \rho_1, \dots, \rho_t \rangle$ is said to be *of the same type as (similar to) \mathcal{A}* if $s = t$ and for every i the relation ρ_i has the same arity as θ_i . A *homomorphism* from \mathcal{A} to \mathcal{B} is a map $f : A \rightarrow B$ such that $f(\theta_i) \subseteq \rho_i$ for every $i = 1, \dots, s$.

Let \mathcal{A} be finite structure of finite type, and let $CSP(\mathcal{A})$ denote the following decision problem:

Presented by M. Valeriote.

Received January 8, 2005; accepted in final form April 3, 2006.

2000 *Mathematics Subject Classification*: Primary 68Q25; Secondary 08B05.

Key words and phrases: NP-complete, constraint satisfaction problem, bounded width problem, variety, Hobby–McKenzie types, poset.

The first author's research is supported by a grant from NSERC and the Centre de Recherches Mathématiques. The second author's research is supported by OTKA no. 034175 and 48809 and T 037877. Part of this research was conducted while the second author was visiting Concordia University in Montréal and also when the first author was visiting the Bolyai Institute in Szeged. The support of NSERC, OTKA and the Bolyai Institute is gratefully acknowledged.

- $\text{CSP}(\mathcal{A})$

Input: A structure \mathcal{I} of the same type as \mathcal{A} .

Question: Is there a homomorphism from \mathcal{I} to \mathcal{A} ?

Many natural decision problems such as SAT and graph k -colourability can be expressed in this form, and this class of decision problems has generated much attention in the past few years (see for example [5, 6, 7, 9, 10, 11, 13, 17, 20, 21, 22, 23, 25]), namely in connection with a conjecture about its expressive power, first stated by Feder and Vardi in 1993 [16]: *Every constraint satisfaction problem is either in \mathbf{P} or \mathbf{NP} -complete.* An approach that exploits the notions of polymorphisms and finite algebras was initiated by Jeavons [20] and further developed in collaboration with Bulatov and Krokhin [11]. Subsequently, Bulatov managed to prove the conjecture in various special cases [6], [7].

Feder and Vardi [17] studied two special types of CSP's they termed *problems of bounded width* and those with the *ability to count*. They argued that all CSP's should be, up to polynomial-time equivalence, in one of these classes, and that the second class should satisfy the dichotomy conjecture. Problems of bounded width are defined in terms of the language Datalog, or equivalently via certain two-player games (see 2.2 below): they are polynomial-time solvable by local consistency algorithms (see for instance [23], [13], [14], [21]).

In this paper, we investigate various algebraic properties of the CSP's associated to problems of bounded width. We show that the algebras associated to bounded width CSP's form a class that is "robust" under various constructions (Lemma 3.3). We give a criterion, in terms of the Hobby–McKenzie types [19], to determine if an algebra is not of bounded width (Theorem 4.2). We use this result to show that various graph-related decision problems do not have bounded width (Proposition 5.2), answering a question posed by Nešetřil and Zhu [26]; the reader is referred to the recent [1] for related results.

In [17], Feder and Vardi associate with every finite structure a finite poset we will call the *Feder–Vardi poset* of the structure. The *retraction problem* for a finite poset P is the following decision problem:

- $\text{Ret}(P)$

Input: A finite poset I containing P .

Question: Is there a retraction from I onto P ?

Feder and Vardi proved that the CSP for the structure and the retraction problem for its associated poset are polynomial-time equivalent. In light of this result, a dichotomy for the retraction problems of finite posets yields the dichotomy for the CSPs of arbitrary structures.

For a finite poset $P = \{p_1, \dots, p_n\}$, let \mathcal{P} denote the relational structure $\mathcal{P} = \langle P; \leq, \{p_1\}, \dots, \{p_n\} \rangle$. It is well known that for any finite poset P , $\text{Ret}(P)$ and $\text{CSP}(\mathcal{P})$ are polynomial-time equivalent (see for instance [25]). It is natural to define the *width* of $\text{Ret}(P)$ as the width of $\text{CSP}(\mathcal{P})$. In this paper we prove that if P is the Feder–Vardi poset of a structure \mathcal{A} with a single relation, and $\text{Ret}(P)$ has bounded width then $\text{CSP}(\mathcal{A})$ also has bounded width (Corollary 5.4). We use this result to present a finite poset P whose associated order-primal algebra generates a variety that omits type 1 but admits type 2, provided $\mathbf{P} \neq \mathbf{NP}$ (Proposition 5.5).

2. Problems of bounded width

For ease of notation, we shall feel free in the sequel to denote relational structures by their underlying universe, e.g. we denote the structure \mathcal{A} simply by A .

2.1. An algorithm. Let I be a relational structure. As usual, if K is a non-empty subset of I , the *substructure induced by K* is the structure with universe K and whose relations are those of I restricted to K . Given similar structures I and A and a subset K of I , we let $\text{Hom}(K, A)$ denote the set of homomorphisms from K to A where K is viewed as a substructure of I . Let k be a positive integer. We call the subsets of size at most k of a set I the *k -subsets* of I .

Fix a structure A and integers $1 \leq l < k$. We now describe an algorithm:

(l, k) -algorithm

Input: A structure I similar to A .

Initial step: To every k -subset K of I assign a relation ρ_K that consists of all maps in $\text{Hom}(K, A)$ viewed as $|K|$ -tuples;

Iteration step: Choose, if they exist, two k -subsets H and K of I with $|H \cap K| \leq l$ such that there is a map in ρ_H whose restriction to $H \cap K$ is not equal to the restriction to $H \cap K$ of any map in ρ_K , and throw out from ρ_H all such maps. If no such H and K are found then stop and output the current relations assigned to the k -subsets of I .

The relations given in the initial step are called the *input relations* of the (l, k) -algorithm. We refer to the relations ρ_K obtained during the algorithm as *k -relations*. The k -relations obtained at the end of the algorithm are called the *output relations*. Since the number of k -subsets of I is $\mathcal{O}(n^k)$ where n is the size of the instance, and in each iteration step the sum of the sizes of the k -relations is decreasing, the algorithm stops in polynomial time in the size of the structure I .

Notice that the choice of the pair H and K in each iteration step of the algorithm is arbitrary. So the (l, k) -algorithm has several different versions depending on the

method of the choice of the pair H and K . By using induction we prove that the output relations produced by the (l, k) -algorithm are the same for all versions of the (l, k) -algorithm. The induction argument is based on the obvious fact that if I and J are two structures with the same input relations then by running the (l, k) -algorithm for both structures in the same way the output relations will be the same.

Proposition 2.1. *Let A and I be similar relational structures. Then any two versions of the (l, k) -algorithm for I output the same relations.*

Proof. We use induction on the sum of the sizes of the input relations. If the sum is 0 there is nothing to prove. We are also done if in the first iteration step there exist no sets H and K with the required property. So let us assume that a version of the algorithm starts by choosing k -subsets H and K and another one starts by choosing H' and K' . Now, by the induction hypothesis all versions of the (l, k) -algorithm where we fix the choice H and K in the first iteration step yield the same output relations. This is also true with H' and K' instead of H and K . We show that there is a version of the (l, k) -algorithm that starts with choosing H and K , and another that starts with choosing H' and K' so that at one point they yield the same k -relations for I . This will conclude the proof.

Up to symmetry we have the following possibilities for H, K, H' and K' :

- (1) H, K, H' and K' are all different: in the second step of the first version we choose H' and K' and in the second step of the second version we choose H and K . Then at the end of the second step both versions yield the same k -relations.
- (2) $H = H'$ and $K = K'$: the two versions yield the same k -relations at the end of the first step.
- (3) $H = K'$ and $K = H'$: in the second step of the first version we choose K and H and in the second step of the second version we choose H and K . Then at the end of the second step the k -relations given by the two versions are the same.
- (4) $H = H'$ and H, K, K' are all different: it might happen that the k -relations given by the two versions are the same at the end of the first step. If not then in the second step of the first version we choose H and K' and in the second step of the second version we choose H and K . One of the choices might not be possible in the second iteration step; in that case we stay with the first step of the corresponding version. So either at the end of the first step of one of the versions and at the end of the second step of the other version or at the end of the second step of both versions the corresponding k -relations are the same.
- (5) $H = K'$ and H, K, H' are all different: in the second step of the first version we choose H' and H . In the second step of the second version we choose H and K . If the corresponding k -relations after two steps are not the same for

both versions then for the second version in the third step we choose H' and H , which is now possible. So the first version at the end of the second step yields the same k -relations as the second version at the end of the third step.

- (6) $K = K'$ and H, K, H' are all different: in the second step of the first version we choose H' and K and in the second step of the second version we choose H and K . Then at the end of the second step both versions yield the same k -relations. \square

2.2. A two-player game. Let I be a relational structure similar to A and let $1 \leq l < k$ be integers. We present a two-player combinatorial game as in [17] (see also [23] and [13]), the (l, k) -game on I : first, Player 1 (the *Spoiler*) selects a k -subset K of I ; Player 2 (the *Duplicator*) then picks a map $f: K \rightarrow A$. Then the Spoiler selects another k -subset K' such that $|K \cap K'| \leq l$; the Duplicator now picks a map $f': K' \rightarrow A$ such that $f|_{K \cap K'} = f'|_{K \cap K'}$. Proceeding in this fashion, the Spoiler wins if at some point the map picked by the Duplicator is not a homomorphism from its domain considered as a substructure of I to A . As usual, we say that the Spoiler has a *winning strategy on I* if the Spoiler can play so that the Duplicator, whatever sequence of moves is chosen, is eventually forced to pick a map which is not a homomorphism.

The notions of (l, k) -game and (l, k) -algorithm are connected by the following proposition:

Proposition 2.2. *Let A and I be similar relational structures. Then the (l, k) -algorithm for I yields empty output relations if and only if the Spoiler has a winning strategy in the (l, k) -game for I .*

Proof. Notice first that by definition of the algorithm, one of the output relations is empty if and only if all the output relations are empty. If the output relations given by the (l, k) -algorithm are nonempty then in the (l, k) -game the Duplicator always picks an appropriate map from the output relation assigned to the k -subset that was selected by the Spoiler in the prior step, preventing the Spoiler from winning the game. Suppose now that the output relations of the (l, k) -algorithm are empty. By proceeding backwards in the course of the algorithm, the Spoiler picks the k -subsets $K_i, i = 1, \dots, t$, as follows: the first choice K_1 of the Spoiler is the k -subset whose k -relation becomes empty first when we carry out the (l, k) -algorithm for I . The second choice K_2 of the Spoiler is the k -subset such that K_1 and K_2 were chosen in the iteration step that eliminated the first choice of the Duplicator from the relation ρ_{K_1} . The third choice K_3 is the k -subset such that K_2 and K_3 were chosen in the iteration step that eliminated the second pick of the Duplicator from ρ_{K_2} . We define the remaining K_i in this fashion until we get to the input relation ρ_{K_t} . Observe that this strategy is well defined, since in the i -th step of the game the

Duplicator is forced to pick a map outside of ρ_{K_i} of the appropriate step of the (l, k) -algorithm. So in the t -th step the Duplicator has to pick a map that is not a homomorphism. \square

Clearly, if the output relations of the (l, k) -algorithm for I are empty then there is no homomorphism from I to A ; however, it might be that the converse does not hold. Following [17], we say that a problem $\text{CSP}(A)$ has *width* (l, k) if for any relational structure I for which the Spoiler has no winning strategy in the (l, k) -game, there exists a homomorphism from I to A . We say that $\text{CSP}(A)$ has *width* l if it has width (l, k) for some k , and that $\text{CSP}(A)$ has *bounded width* if it has width l for some l . By the last result, it follows that $\text{CSP}(A)$ has bounded width if for some choice of parameters l and k the (l, k) -algorithm correctly decides the problem $\text{CSP}(A)$: in particular, we get that $\text{CSP}(A) \in \mathbf{P}$. For example, any relational structure of finite type whose relations are invariant under a semilattice operation, or a near-unanimity operation, has bounded width [22], [21]. We shall give several examples of problems that do not have bounded width in section 5.

Let l, k, l', k' be integers such that $1 \leq l < k$ and $1 \leq l' < k'$ with $l' \geq l$ and $k' \geq k$. It can be easily verified that if $\text{CSP}(A)$ has width (l, k) then it has width (l', k') . For convenience, we introduce the following terminology: we will say that a relational structure A is an (l, k) -*structure* if $\text{CSP}(A)$ has width (l, k) . Finally, we will say that a structure I similar to A is (l, k) -*consistent (for A)* if running the (l, k) -algorithm on I yields non-empty output relations. Note that there are related but different notions of consistency in the literature, see for example [13].

3. Algebras of bounded width

For basic results in universal algebra and tame congruence theory we refer to [12], [19] and [28].

A *clone* on a finite set is set of finitary operations containing all projections and closed under composition. To any relational structure A is naturally associated a clone: the *clone of A* is the set of operations preserving all the relations of A . The *relational clone* of A is the set of finitary relations on the base set of A preserved by all operations in the clone of A . If A is a relational structure on the universe of an algebra \mathbb{A} , we say that the algebra \mathbb{A} is an *algebra for A* (or that A is a *relational structure for \mathbb{A}*) if the clone of term operations of \mathbb{A} coincides with the clone of the structure A .

We wish to investigate the properties of the algebras associated to structures of bounded width, which prompts the following natural definition:

Definition 3.1. We say that a finite algebra \mathbb{A} has *bounded width* if for every relational structure \mathcal{A} (of finite type) whose base set coincides with the universe of

\mathbb{A} and whose relations are subalgebras of finite powers of \mathbb{A} , the problem $\text{CSP}(\mathcal{A})$ has bounded width.

We will show in this section that the class of algebras of bounded width is closed under familiar algebraic constructions, such as formation of products, subalgebras and homomorphic images (Lemma 3.3). This is then used to show that bounded width is preserved under interpretability of varieties (Theorem 3.4).

We begin with a result that ensures that if \mathbb{A} is an algebra for an (l, k) -structure then it is an algebra of bounded width. Its proof, although more involved, is similar in flavour to the proof of the following result of Jeavons [20]: *If B is a relational structure of finite type whose base set coincides with the base set of A and whose relations are in the relational clone of A , then $\text{CSP}(B)$ is polynomial-time reducible to $\text{CSP}(A)$.*

Lemma 3.2. *Let A be an (l, k) -structure. If B is a relational structure whose base set coincides with the base set of A and whose relations are in the relational clone of A , then B is an (l', k') -structure for some l' and k' .*

Proof. If B satisfies the conditions of the lemma, then each of its basic relations is obtained from those of A in finitely many steps by using the following constructions [3]:

- (1) removing a relation,
- (2) adding a relation obtained by permuting the variables of a relation,
- (3) adding the intersection of two relations of the same arity,
- (4) adding the product of two relations,
- (5) adding the equality relation,
- (6) adding a relation obtained by projecting an n -ary relation to its first $n - 1$ variables.

It thus suffices to prove that if B is obtained by any of these constructions from A , then B is an (l', k') -structure for some l' and k' . By the remarks at the end of the last section, we may safely assume throughout this proof that $l = k - 1$; we shall also assume, when choosing k' , that $l' = k' - 1$. In each of the 6 cases we shall choose a convenient value of k' and an arbitrary (l', k') -consistent structure I for B and construct from it a relational structure J similar to A . It will turn out that J is an (l, k) -consistent structure for A and by the use of a homomorphism from J to A it will be straightforward to define a homomorphism from I to B . It will follow that B is an (l', k') -structure.

Case 1: Let us suppose first that B is obtained from A by removing a relation of A . Let J be the relational structure obtained from I by supplementing the relations of I by an empty relation corresponding to the relation removed from A . Observe that when we carry out the (l, k) -algorithm for I and J the resulting k -relations

are the same on each k -subset in each step. Hence if I is (l, k) -consistent, so is J . Since A is an (l, k) -structure there is a homomorphism from J to A . Clearly, the same map is a homomorphism from I to B .

Case 2: Let us suppose now that B is obtained from A by adding a relation s via a permutation of the variables of the relation r of A by a permutation π . Let J be the relational structure obtained from I by deleting s_I and replacing r_I by $r_J = r_I \cup s'_I$, where s'_I is obtained by permuting the variables of s_I according to the inverse of π . Observe that the (l, k) -algorithm works in the same way for both I and J . So if I is (l, k) -consistent so is J . Since A is an (l, k) -structure there is a homomorphism from J to A . Clearly, the same map is also a homomorphism from I to B .

Case 3: let r and s be two relations of the same arity of A . Let t denote the intersection of r and s . Let us suppose that B is obtained from A by adding the relation t . Let J be the relational structure obtained from I by deleting t_I and replacing r_I by $r_J = r_I \cup t_I$, and s_I by $s_J = s_I \cup t_I$. Observe that the (l, k) -algorithm works in the same way for both I and J . If I is (l, k) -consistent, so is J . Since A is an (l, k) -structure there is a homomorphism from J to A . Clearly, the same map is also a homomorphism from I to B .

Case 4: let r and s be two relations of A . Let t denote the product of r and s , and let B be obtained from A by adding the relation t . By the remark at the end of section 2 we may assume that k is at least the sum of the arities of r and s . Suppose that I is (l, k) -consistent, and let J be the relational structure obtained from I by deleting t_I and replacing r_I by $r_J = r_I \cup t^1$, and s_I by $s_J = s_I \cup t^2$ where t^1 is the projection of t_I onto the variables of the r -part of t_I and t^2 is the projection of t_I onto the variables of the s -part of t_I . We claim that each relation output by the (l, k) -algorithm on I is contained in the corresponding input relation for J . Indeed, let $f: K \rightarrow A$ be a homomorphism when K is a substructure of I . If it is *not* a homomorphism when K is a substructure of J , it means that there is a tuple $e \in t_I$ whose variables are not all contained in K such that, without loss of generality, $f(e^1)$ is not in r , where e^1 denotes the projection of e on its r -part. Choose a k -subset H that does contain all the variables of e ; this is possible by our choice of k . Then obviously the restriction of f to $H \cap K$ cannot extend to a homomorphism $g: H \rightarrow A$, and hence f is not in the output relation associated to K . From this it follows that J is (l, k) -consistent, and so there is a homomorphism from J to A . Clearly, the same map is also a homomorphism from I to B .

Case 5: Let us suppose now that B is obtained from A by adding the equality relation. We shall prove that B is an (l', k') -structure where $k' = 2k$. Let I be (l', k') -consistent for B . Let θ be the reflexive, symmetric, transitive closure of θ' where θ' is the relation of I corresponding to equality in B . Clearly, θ is an equivalence relation on I . We define a structure J similar to A ; its base set is the

set of θ -blocks, and its relations are defined as follows: for each relation r_I of I distinct from θ' we define a relation r_J on the θ -blocks of I by stipulating that

$$(I_1, \dots, I_n) \in r_J \text{ iff } \exists h_1 \in I_1, \dots, \exists h_n \in I_n \text{ such that } (h_1, \dots, h_n) \in r_I.$$

Let U and V be two $2k$ -subsets of I and let $u_1, \dots, u_k \in U$ and $v_1, \dots, v_k \in V$ such that $u_1\theta v_1, \dots, u_k\theta v_k$. Then there exist a sequence $w_{1,j}, \dots, w_{k,j}$, $j = 0, \dots, t$ of elements for which (i) the pairs $(w_{1,j}, w_{1,j+1}), \dots, (w_{k,j}, w_{k,j+1})$ are contained in the reflexive symmetric closure of θ' , (ii) $w_{1,0} = u_1, \dots, w_{k,0} = u_k$ and (iii) $w_{1,t} = v_1, \dots, w_{k,t} = v_k$. Hence there is a sequence of $2k$ -subsets $U = W_0, W_1, \dots, W_t = V$ such that $\{w_{1,j}, \dots, w_{k,j}, w_{1,j+1}, \dots, w_{k,j+1}\} \subseteq W_j$ for all j with $1 \leq j \leq t-1$. So for any map f in the output relation ρ_U there is a map g in the output relation ρ_V such that $f(u_1) = g(v_1), \dots, f(u_k) = g(v_k)$.

We fix a set of representatives of the θ -blocks of I . Let H be any k -subset of J . We prove that the output relation μ_H given by the (l, k) -algorithm for J is nonempty. Let U be any $2k$ -subset of I containing the representatives from all θ -blocks of H . Let f be any map in the output relation ρ_U (there exists such a map since I is (l', k') -consistent). Define a map $f': H \rightarrow A$ by setting $f'(T) = f(t)$ where t is the fixed representative in the θ -block T . First we show that f' is a homomorphism. Indeed, if $(T_1, \dots, T_n) \in r_J$ where $T_1, \dots, T_n \in H$, then there exist $v_1 \in T_1, \dots, v_n \in T_n$ such that $(v_1, \dots, v_n) \in r_I$. Let t_1, \dots, t_n be the fixed representatives of the θ -blocks T_1, \dots, T_n respectively. By the preceding paragraph, for any $2k$ -subset V of I that contains v_1, \dots, v_n there is a g in the output relation ρ_V such that $f(t_1) = g(v_1), \dots, f(t_n) = g(v_n)$. Since g is in an output relation, g preserves r , and hence $(f'(T_1), \dots, f'(T_n)) = (f(t_1), \dots, f(t_n)) \in r$. Thus, for any $2k$ -subset U containing the fixed representatives of the blocks of H and any $f \in \rho_U$, the corresponding map $f': H \rightarrow A$ is a homomorphism and so it is in the input relation on H .

By induction we show that the (l, k) -algorithm for J eliminates no f' of the above form from the k -relation on H . Suppose that $H = \{T_1, \dots, T_k\}$ and t_1, \dots, t_k are the fixed representatives of the blocks T_1, \dots, T_k , respectively. Let f' be any map in μ_H constructed as above, i.e., such that there is an f from the output relation ρ_U of the (l', k') -algorithm for I where U is any $2k$ -subset of I which contains $\{t_1, \dots, t_k\}$ and $f'(T_1) = f(t_1), \dots, f'(T_n) = f(t_n)$. There is only one way for f' to be removed from μ_H : in some iteration step of the (l, k) -algorithm for J we choose the k -subsets H and some K . Then there is a $2k$ -subset V of I such that V contains all fixed representatives from the θ -blocks of $H \cup K$. Since f is in ρ_U there is a map g in ρ_V such that $f(t_1) = g(t_1), \dots, f(t_k) = g(t_k)$. Then by the induction hypothesis $g' \in \mu_K$. So f' will not be eliminated from μ_H . Thus, J is (l, k) -consistent.

Since A is an (l, k) -structure there is a homomorphism from J to A . A homomorphism from J to A composed with the natural map corresponding to θ is a homomorphism from I to B .

Case 6: Let us suppose that B is obtained from A by adding the projection s of an n -ary relation r of A to its first $n - 1$ variables. In this case we want to prove that B is an (l', k') -structure with $k' = k^2$. So let I be any (l', k') -consistent structure for B . We define a structure J similar to A : its base set is the base set of I extended by a new element for each $(n - 1)$ -tuple in s_I ; the relations of J are those of I , except that s_I is removed and r_I is replaced by $r_J = r_I \cup s'_I$, where s'_I is obtained from s_I by extending every $(n - 1)$ -tuple of s_I with the corresponding new element in the base set of J . We shall prove that J is an (l, k) -consistent structure for A .

Let H be any k -subset of J . We choose a k^2 -subset H' of I that contains $I \cap H$ and, additionally, for each element h in $H \setminus I$, all components of the tuple in I that corresponds to h in the definition of J , provided that the number of distinct components in this tuple is at most k . Since H contains at most k elements it is easy to see that such a k^2 -subset exists. Let f' be the restriction of any map f from the output relation on H' to $H \cap I$, and let $f'' : H \rightarrow A$ be any homomorphism extension of f' to H : since I is an (l', k') -consistent structure there exists such an f , and then by the definition of H' its restriction f' is easily seen to extend to a homomorphism f'' on H .

Now, we claim that every $f'' : H \rightarrow A$ defined in this way is contained in ρ_H at each step of the (l, k) -algorithm run for J . We prove the claim by induction. It is certainly true in the initial step. Let us consider an arbitrary iteration step. It suffices to show that f'' is not eliminated from ρ_H in this step. By induction hypothesis each k -relation produced by the algorithm in the previous step contains all maps defined according to the preceding paragraph. Suppose that in our iteration step we choose the k -subsets H and K . Let K' be any k^2 -subset of I defined similarly from K as H' was from H above. Let f'' be obtained from a map f in the output relation for H' as described. Since I is (l', k') -consistent, there is map g in the output relation on K' such that $f|_{H' \cap K'} = g|_{H' \cap K'}$. Let $g'' : K \rightarrow A$ be the map obtained from g in the same way as f'' was obtained from f , and furthermore by using the same extension on $(H \cap K) \setminus I$ that we used in defining f ; by definition of J this is clearly feasible. Then by induction hypothesis, $g'' \in \rho_K$ and $f''|_{H \cap K} = g''|_{H \cap K}$. So after carrying out the iteration step, f'' remains in ρ_H , which proves the claim. This means that J is an (l, k) -consistent structure for A . Hence J admits a homomorphism to A . Clearly, the restriction of this map to I is a homomorphism from I to B . \square

A *variety* is a class of algebras of the same type closed under formation of subalgebras, homomorphic images and products. For any algebra \mathbb{A} there is a smallest

variety containing \mathbb{A} , denoted by $\mathcal{V}(\mathbb{A})$ and called the *variety generated by \mathbb{A}* . It is well known that any variety is generated by an algebra and that any member of $\mathcal{V}(\mathbb{A})$ is a homomorphic image of a subalgebra of a power of \mathbb{A} .

A finite algebra \mathbb{A} is called *locally tractable* if the problem $\text{CSP}(\mathcal{A})$ is in \mathbf{P} for every relational structure \mathcal{A} of finite type whose base set coincides with the universe of \mathbb{A} and whose relations are subalgebras of finite powers of \mathbb{A} . It follows from results in [11] and [9] that *if a finite algebra \mathbb{A} is locally tractable, then so is every finite algebra in $\mathcal{V}(\mathbb{A})$* . An analogous statement is valid for bounded width algebras:

Lemma 3.3. *Every finite algebra in the variety generated by a bounded width algebra has bounded width.*

Proof. Let \mathbb{A} be a bounded width algebra. It suffices to show that every subalgebra, homomorphic image and finite power of \mathbb{A} has bounded width. Let \mathbb{B} be a subalgebra, a homomorphic image, or a finite power of \mathbb{A} and let B be a relational structure on the base set of \mathbb{B} such that the relations of B are subalgebras of finite powers of \mathbb{B} . We shall prove that B is an (l, k) -structure for some l and k .

Suppose first that \mathbb{B} is a subalgebra of \mathbb{A} . Let A be the relational structure whose base set is the base set of \mathbb{A} and whose relations are all the relations of B and the base set of B as a unary relation. Notice that the relations of A are subalgebras of finite powers of \mathbb{A} . Since \mathbb{A} has bounded width, A is an (l, k) -structure for some l and k . As in the proof of the last lemma, we assume for simplicity that $l = k - 1$. Now, given any (l, k) -consistent structure I for B we define a structure J for A as follows: we take J to be I with all of its relations adding the base set of I as a unary relation. Observe that when we carry out the (l, k) -algorithm for I and J in the initial steps the input relations defined on each k -subset agree. Hence they will agree in each step of the algorithm. Since I is (l, k) -consistent, J is also (l, k) -consistent. Since A is an (l, k) -structure, there is a homomorphism from J to A . Clearly, the same map is a homomorphism from I to B .

Secondly, suppose that \mathbb{B} is a homomorphic image of \mathbb{A} under the homomorphism h . This time let A be the relational structure whose base set is the universe of \mathbb{A} and whose relations are the preimages under the homomorphism h of all the relations of B . Notice that the relations of A are subalgebras of the finite powers of \mathbb{A} . Hence A is an (l, k) -structure for some l and k . We show that any (l, k) -consistent structure I for B is also an (l, k) -consistent structure for A . Observe that each k -relation of the (l, k) -algorithm carried out on I for A contains the preimage of the corresponding output relation of the (l, k) -algorithm carried out on I for B . Hence, since I is (l, k) -consistent for B , I is also (l, k) -consistent for A . Since A is an (l, k) -structure, there is a homomorphism f from I to A . Clearly, the map hf is a homomorphism from I to B .

Finally suppose that $\mathbb{B} = \mathbb{A}^n$. Let A be the relational structure on the base set of \mathbb{A} , with the following relations. If r is an s -ary relation of B , define r_0 to be the sn -ary relation such that, if $(b_1, \dots, b_s) \in r$ with $b_i = (a_{1,i}, \dots, a_{n,i})$, we put the sn -tuple $(a_{1,1}, \dots, a_{1,s}, \dots, a_{n,1}, \dots, a_{n,s})$ in r_0 . Note that the sn -ary relations obtained in this way are subalgebras of finite powers of \mathbb{A} . Hence A is an (l, k) -structure for some l and k . For any (l, k) -consistent structure I for B we define a structure J for A as follows. We take J to be the union of n disjoint copies of I with one sn -ary relation for each s -ary relation of I . An sn -tuple in the new relation on J is formed by the n copies of an s -tuple in the old relation on I , that is, if (b_1, \dots, b_s) is in the old relation on I and $b_{i,j}$ is the i -th copy of b_j in J , $i = 1, \dots, n$, $j = 1, \dots, s$, then $(b_{1,1}, \dots, b_{1,s}, \dots, b_{n,1}, \dots, b_{n,s})$ is in the new relation.

We show that J is (l, k) -consistent for A . Choose a k -subset H_0 in J . The set H_0 decomposes into n subsets, D_i , $i = 1, \dots, n$, according to the n copies of I in J . Then there is a k -subset H in I whose multiple copies in J contain all elements of H_0 . Let t be any element in the output relation on H when the (l, k) -algorithm is carried out on I . Let t_0 be the tuple naturally obtained from t , considered as a map from H_0 to A , that is, for every $d \in D_i$, $t_0(d) = t(h)$ where $h \in H$ and d is the i -th copy of h in J .

It is easy to see by induction that t_0 is in the output relation related to H_0 . Since I is (l, k) -consistent, J is also (l, k) -consistent. Because A is an (l, k) -structure we get that there is a homomorphism f from J to A . Let $g: I \rightarrow B$, $x \mapsto (f(x_1), \dots, f(x_n))$ be the map where x_i is the i -th copy of x in J . Clearly, g is a homomorphism from I to B . \square

We conclude this section with a result that states that bounded width is preserved under interpretation of varieties; this will be used in the next section to produce a criterion to prove that many CSP's do not have bounded width. For the purposes of Theorem 3.4 we shall define interpretability as follows: if \mathbb{B} is an algebra, we say that a variety \mathcal{V} *interprets in* $\mathcal{V}(\mathbb{B})$ if there is an algebra in \mathcal{V} with the same universe as \mathbb{B} , all of whose term operations are term operations of \mathbb{B} . We shall also require the following alternative definition in Theorem 5.3: the variety $\mathcal{V}(\mathbb{A})$ interprets in the variety $\mathcal{V}(\mathbb{B})$ if and only if there exists a clone homomorphism from the clone of term operations of \mathbb{A} to the clone of term operations of \mathbb{B} , where a map between clones is called a *clone homomorphism* if it preserves arity, maps projections to projections and commutes with composition (see [19] page 131 for details).

Theorem 3.4. *If \mathbb{A} and \mathbb{B} are finite algebras such that $\mathcal{V}(\mathbb{A})$ interprets in $\mathcal{V}(\mathbb{B})$ and \mathbb{A} has bounded width, then \mathbb{B} also has bounded width.*

Proof. Let B be a finite relational structure of finite type on the base set of \mathbb{B} such that the relations of B are all subalgebras of finite powers of \mathbb{B} . It suffices to show

that B is an (l, k) -structure for some l and k . Since the variety $\mathcal{V}(\mathbb{A})$ interprets in the variety $\mathcal{V}(\mathbb{B})$, $\mathcal{V}(\mathbb{A})$ contains an algebra \mathbb{C} whose universe is that of \mathbb{B} and whose clone of term operations is contained in that of \mathbb{B} . So the base set of B coincides with the base set of \mathbb{C} and the relations of B are subalgebras of finite powers of \mathbb{C} . Since \mathbb{A} has bounded width, by Lemma 3.3, $\mathbb{C} \in \mathcal{V}(\mathbb{A})$ also has bounded width. Hence B is an (l, k) -structure for some l and k . \square

4. An omitting-type criterion for bounded width

In this section we present a criterion to determine if certain algebras are not of bounded width, based on the notion of the *type set* of an algebra and of a variety; the reader should consult [19] or [12] for details. Roughly speaking, the type set of a finite algebra is a subset of the set $\{1, 2, 3, 4, 5\}$ whose elements are called *types* and correspond to certain classes of algebras: 1 to unary algebras, 2 to vector spaces over finite fields, 3 to Boolean algebras, 4 to distributive lattices and 5 to semilattices. If \mathcal{V} is a variety, the *type set of \mathcal{V}* is the union of the type sets of the finite algebras in \mathcal{V} . We say that an algebra or variety *admits (omits) type i* when i is (is not) in its type set. Just to have a feeling how types affect the shape of algebras of a variety we mention that if for instance a variety \mathcal{V} omits types 1 and 2 one can think of the members of \mathcal{V} as certain amalgams of Boolean algebras, distributive lattices and semilattices. A variety \mathcal{V} is called *locally finite* if all finitely generated members of \mathcal{V} are finite. It is easy to verify that any variety generated by a finite algebra is locally finite. In what follows, we shall require characterisations of certain of these varieties that omit some of the five types; these so-called *omitting-types theorems* can be found in Chapter 9 of [19].

The following connection between the type set of the variety generated by the algebra associated to a structure and the complexity of the associated CSP is a consequence of results in [11]: *let \mathbb{A} be a finite, idempotent algebra such that $\mathcal{V}(\mathbb{A})$ admits type 1. If A is a structure for \mathbb{A} , then the problem $\text{CSP}(A)$ is **NP**-complete.*

We shall now use Theorem 3.4 to prove a parallel result, namely that the variety generated by an idempotent algebra of bounded width must omit types 1 and 2 (Theorem 4.2) (we will end the section by proving that there is no loss of generality in considering idempotent algebras (Lemma 4.3)). For this we need a lemma whose proof can be put together from results contained in [19] and [29]. A congruence θ of an algebra \mathbb{A} is *Abelian* if for any n -ary polynomial f of \mathbb{A} and any $u, v, x_1, \dots, x_n, y_1, \dots, y_n \in A$, if $u\theta v$ and $x_i\theta y_i$ for all $i = 1, \dots, n$, then $f(u, x_1, \dots, x_n) = f(u, y_1, \dots, y_n)$ implies $f(v, x_1, \dots, x_n) = f(v, y_1, \dots, y_n)$. An algebra is *Abelian* if all its congruences are Abelian. An algebra is *affine* if its clone of polynomial operations coincides with the clone of polynomial operations of a

module on the same universe. Affine algebras are prototypical examples of Abelian algebras.

Lemma 4.1. *For a locally finite idempotent variety \mathcal{V} the following are equivalent:*

- (1) \mathcal{V} omits types 1 and 2.
- (2) The only Abelian congruence of any algebra in \mathcal{V} is the identity relation.
- (3) \mathcal{V} does not interpret in any variety generated by an affine algebra.

Proof. The equivalence of (1) and (2) is given by Theorem 9.10 in [19], and (2) implies (3) is trivial. We prove that (3) implies (2). If (2) does not hold, then \mathcal{V} contains an algebra \mathbb{A} that has a nontrivial Abelian congruence, i.e., some block of this congruence contains at least two elements, call them a and b . Observe that if $\gamma \leq \theta$ are congruences of an algebra and θ is Abelian, then γ is also Abelian. Moreover, the restriction of Abelian congruences to subalgebras are Abelian congruences of the respective subalgebras. Finally, since the variety \mathcal{V} is idempotent, every block of a congruence is a subalgebra. Hence, by taking the subalgebra of \mathbb{A} generated by $\{a, b\}$ and using the local finiteness of \mathcal{V} , we may assume that the algebra \mathbb{A} is finite, and in fact, by choosing such an algebra with minimum cardinality we may also assume that it has no non-trivial subalgebras. Hence \mathbb{A} is a simple Abelian algebra. By [29], the only Abelian idempotent algebras with no non-trivial subalgebras are affine algebras and algebras with two elements and trivial clone. In any case, \mathbb{A} is an algebra whose clone is contained in the clone of an affine algebra, and so \mathcal{V} interprets in a variety generated by an affine algebra. \square

Theorem 4.2. *If \mathbb{A} is a finite idempotent algebra of bounded width, then $\mathcal{V}(\mathbb{A})$ omits types 1 and 2.*

Proof. Let \mathbb{A} be any finite idempotent algebra such that $\mathcal{V}(\mathbb{A})$ admits type 1 or 2. Then by Lemma 4.1, $\mathcal{V}(\mathbb{A})$ interprets in the variety generated by an affine algebra \mathbb{C} . In fact, since it is idempotent, $\mathcal{V}(\mathbb{A})$ interprets in $\mathcal{V}(\mathbb{B})$ where \mathbb{B} is an algebra whose base set is C and whose clone consists of all idempotent operations of \mathbb{C} . Let us consider the structure $D = \langle B; \{0\}, \{(x, y, z) : x + y + z = \alpha\} \rangle$, where B is the base set of \mathbb{B} , and α is a fixed non-zero element of B . The relations of D are preserved by all operations of \mathbb{B} and D is a structure which has the ability to count (see [17] page 85). So by Theorem 31 in [17], D is not an (l, k) -structure for any l and k . Hence \mathbb{B} does not have bounded width. Now, Theorem 3.4 implies that \mathbb{A} does not have bounded width either. \square

The preceding criterion may be used to identify algebras that do not have bounded width even if they are not idempotent, as the next lemma shows. Given an algebra \mathbb{A} and a subset B of its universe, let $\mathbb{A}|_B$ denote the algebra with universe B whose basic operations are the restriction to B of every term operation

of \mathbb{A} that preserves B . The *idempotent reduct* of an algebra \mathbb{A} is the algebra on the same universe whose term operations are the idempotent term operations of \mathbb{A} . We remark that some of the constructions in the proof of the lemma are similar to those in Theorems 3.3 and 3.7 of [11].

Lemma 4.3.

- (1) *Let \mathbb{A} be a finite algebra, and let r be a unary term of \mathbb{A} such that $r^2 = r$. Let $B = r(A)$. Then the algebra $\mathbb{A}|_B$ has bounded width if and only if \mathbb{A} has bounded width.*
- (2) *Let \mathbb{A} be a surjective algebra and let \mathbb{B} be its idempotent reduct. Then \mathbb{A} has bounded width if and only if \mathbb{B} has bounded width.*

Proof. (1) (\Leftarrow) Assume that \mathbb{A} has bounded width. Let $\mathbb{B} = \mathbb{A}|_B$, and let $\mathcal{B} = \langle B; \theta_1, \dots, \theta_s \rangle$ be a structure where each θ_i is a subuniverse of a finite power of \mathbb{B} . Consider the structure $\mathcal{A} = \langle A; \theta'_1, \dots, \theta'_s \rangle$ where θ'_i is the subalgebra generated by θ of the appropriate power of \mathbb{A} . Then \mathcal{A} is an (l, k) -structure for some l and k . As before we assume that $l = k - 1$. Let I be an (l, k) -consistent instance for \mathcal{B} ; we must show that there is a homomorphism from I to \mathcal{B} . Since $\theta_i \subseteq \theta'_i$ for each i , it is clear that I is an (l, k) -consistent instance for \mathcal{A} , and hence there is a homomorphism $\phi: I \rightarrow \mathcal{A}$; to finish the proof it suffice to prove that $r \circ \phi$ is a homomorphism from I to \mathcal{B} . It is clear that $r(\theta'_i) \subseteq \theta'_i \cap B^k$, where k is the arity of θ_i ; we claim that this last set is contained in θ_i . Indeed, if $\bar{x} \in \theta'_i \cap B^k$, then it is of the form $\bar{x} = rg(\bar{y}_1, \dots, \bar{y}_n)$ for some term g of \mathbb{A} and some $\bar{y}_j \in \theta_i$, and clearly $rg|_B$ is a term of \mathbb{B} , so $\bar{x} \in \theta_i$.

(\Rightarrow) Suppose that $\mathbb{B} = \mathbb{A}|_B$ has bounded width, and let $\mathcal{A} = \langle A; \theta_1, \dots, \theta_s \rangle$ be a structure where each θ_i is a subuniverse of a finite power of \mathbb{A} . Consider the structure $\mathcal{B} = \langle B; r(\theta_1), \dots, r(\theta_s) \rangle$. It is easy to see that each $r(\theta_i)$ is a subuniverse of a power of \mathbb{B} . Hence \mathcal{B} is an (l, k) -structure for some l and k . Let I be an (l, k) -consistent instance for \mathcal{A} ; we must show that there is a homomorphism from I to \mathcal{A} . Viewing I as an instance for \mathcal{B} , notice that for any output relation ρ_H of the algorithm for \mathcal{A} , we will have $r(\rho_H)$ contained in the output relation of the algorithm for \mathcal{B} . Thus I is (l, k) -consistent and there exists a homomorphism from I to \mathcal{B} . Since $r(\theta_i) \subseteq \theta_i$ for all i , this is also a homomorphism from I to \mathcal{A} .

(2) (\Leftarrow) Clearly $\mathcal{V}(\mathbb{B})$ interprets in $\mathcal{V}(\mathbb{A})$, so by Theorem 3.4, if \mathbb{B} has bounded width so does \mathbb{A} .

(\Rightarrow) Assume that \mathbb{A} has bounded width, and suppose for convenience that $A = \{1, 2, \dots, n\}$. Let $\mathcal{B} = \langle A; \theta_1, \dots, \theta_s \rangle$ be a structure where each θ_i is a subuniverse of a finite power of \mathbb{B} . Since \mathbb{B} is the idempotent reduct of \mathbb{A} , we may find for each $i = 1, \dots, s$ a subuniverse θ'_i of a finite power of \mathbb{A} and elements $a_1^i, \dots, a_{r_i}^i \in A$ such that

$$(x_1, \dots, x_k) \in \theta_i \iff (x_1, \dots, x_k, a_1^i, \dots, a_{r_i}^i) \in \theta'_i.$$

Define the relation

$$\alpha_{\mathbb{A}} = \{(\sigma(1), \dots, \sigma(n)) : \sigma \in Clo_1(\mathbb{A})\},$$

where $Clo_1(\mathbb{A})$ denotes the set of unary terms of \mathbb{A} . It is immediate to verify that this is a subalgebra of \mathbb{A}^n . It follows that the structure $\mathcal{A} = \langle A; \theta'_1, \dots, \theta'_s, \alpha_{\mathbb{A}} \rangle$ is an (l, k) -structure for some l and k (assume again that $l = k - 1$). Let $\mathcal{I} = \langle I; \mu_1, \dots, \mu_s \rangle$ be an (l, k) -consistent instance for \mathcal{B} ; we must show that there is a homomorphism from \mathcal{I} to \mathcal{B} . We construct an instance $\mathcal{J} = \langle J; \nu_1, \dots, \nu_s, \gamma \rangle$ for \mathcal{A} : its base set J is the disjoint union of the base set of I and of A ; for $i = 1, \dots, s$, the relation ν_i consists of all tuples $(x_1, \dots, x_k, a_1^i, \dots, a_{r_i}^i)$ with $(x_1, \dots, x_k) \in \mu_i$; and finally, $\gamma = \{(1, 2, \dots, n)\}$. We claim that \mathcal{J} is (l, k) -consistent for \mathcal{A} . Indeed, let H be a k -subset of J . Let f be in the relation $\rho_{H \cap I}$ output by the (l, k) -algorithm on \mathcal{I} , and define a tuple g by $g(x) = f(x)$ for $x \in I$ and $g(x) = x$ for $x \in A$. It is easy to verify that such a tuple must be in the relation ρ'_H output by the (l, k) -algorithm on \mathcal{J} . Since \mathcal{I} is (l, k) -consistent, \mathcal{J} is also (l, k) -consistent. Hence there exists a homomorphism $\phi: \mathcal{J} \rightarrow \mathcal{A}$. In particular the restriction of ϕ to A is a term of \mathbb{A} , and since \mathbb{A} is surjective, it is an automorphism, call it σ . We claim that the restriction of $\sigma^{-1}\phi$ to I is a homomorphism from \mathcal{I} to \mathcal{B} . Indeed,

$$\begin{aligned} (x_1, \dots, x_k) &\in \mu_i \\ \implies (x_1, \dots, x_k, a_1^i, \dots, a_{r_i}^i) &\in \nu_i \\ \implies (\phi(x_1), \dots, \phi(x_k), \phi(a_1^i), \dots, \phi(a_{r_i}^i)) &\in \theta'_i \\ \implies (\sigma^{-1}\phi(x_1), \dots, \sigma^{-1}\phi(x_k), \sigma^{-1}\phi(a_1^i), \dots, \sigma^{-1}\phi(a_{r_i}^i)) &\in \theta'_i \\ \implies (\sigma^{-1}\phi(x_1), \dots, \sigma^{-1}\phi(x_k), a_1^i, \dots, a_{r_i}^i) &\in \theta'_i \\ \implies (\sigma^{-1}\phi(x_1), \dots, \sigma^{-1}\phi(x_k)) &\in \theta_i. \end{aligned}$$

□

5. Applications

5.1. Graph-related decision problems. We shall now use the results of the last section to prove that certain graph homomorphism problems cannot be solved using local consistency algorithms. We consider graphs without loops, i.e., relational structures H with a single, binary relation θ which is irreflexive, i.e., $(h, h) \notin \theta$ for all $h \in H$. The graph H is *symmetric* if θ is symmetric, i.e., if $(h, h') \in \theta$ implies $(h', h) \in \theta$, and it is an *oriented graph* if the relation θ is *asymmetric*, i.e., $(h, h') \in \theta$ implies $(h', h) \notin \theta$ (in other words, H is a digraph with no symmetric edges). Given a graph H , the *homomorphism problem for H* is the problem $\text{Hom}(H)$ of deciding

whether there exists a homomorphism from a given graph G to H . It is easy to see that this problem is polynomial-time equivalent to $\text{CSP}(H)$.

In [18], Hell and Nešetřil prove the following: *for a symmetric graph H , the problem $\text{Hom}(H)$ is in \mathbf{P} if H is bipartite, and otherwise it is \mathbf{NP} -complete.* We briefly outline how from a careful inspection of their proof one can extract slightly more (see also [8]).

An n -ary idempotent operation on a set A is a *Taylor operation* if it satisfies, for every $1 \leq i \leq n$, an identity of the form

$$f(x_1, \dots, x_{i-1}, x, x_{i+1}, \dots, x_n) \approx f(y_1, \dots, y_{i-1}, y, y_{i+1}, \dots, y_n)$$

where $x_j, y_j \in \{x, y\}$ for all $1 \leq j \leq n$. It is known that for an idempotent algebra \mathbb{A} , $\mathcal{V}(\mathbb{A})$ omits type 1 if and only if the algebra \mathbb{A} has a Taylor term (Lemma 9.4 and Theorem 9.6 of [19]).

Let \mathbb{A} be an algebra for H , where H is non-bipartite. Let r be a retraction of H onto its core H' , and let \mathbb{B} denote the idempotent reduct of $\mathbb{A}|_{H'}$. To prove that $\text{CSP}(H)$ does not have bounded width, it will suffice, by the above discussion, Theorem 4.2 and Lemma 4.3, to prove that \mathbb{B} does not have a Taylor operation. We shall use the following fact that the reader may easily verify: *if a structure C admits a Taylor operation t , and r is a retraction of C onto a substructure C' , then C' admits the Taylor operation $rt|_{C'}$.*

Each relational structure defined by the indicator and subindicator constructions in the Hell and Nešetřil proof has a base set equal to the universe of an algebra \mathbb{C} in $\mathcal{V}(\mathbb{B})$ and relations that are subalgebras of finite powers of \mathbb{C} . Hell and Nešetřil show that one of these structures retracts to the triangle. Since the triangle admits no Taylor operation, the algebra \mathbb{C} does not admit a Taylor operation, and hence neither does \mathbb{B} .

In [26], Nešetřil and Zhu give a different proof that $\text{CSP}(H)$ does not have bounded width when H is non-bipartite. They also ask whether there exists a direct proof, without assuming that $\mathbf{P} \neq \mathbf{NP}$, that $\text{CSP}(H)$ does not have bounded width when H is a *oriented cycle* and $\text{Hom}(H)$ is \mathbf{NP} -complete. We use algebraic techniques to answer Nešetřil and Zhu's question. Our result relies on a proof of Feder that for a oriented cycle H the problem $\text{CSP}(H)$ is either in \mathbf{P} or \mathbf{NP} -complete [15].

An irreflexive, oriented graph H is a *path* if its vertices can be ordered $\{1, 2, \dots, n\}$ in such a way that, for each $i = 1, \dots, n-1$ exactly one of each pair $\{(i, i+1), (i+1, i)\}$ is an edge, and there are no other edges. An irreflexive oriented graph H with at least 3 vertices is a *cycle* if the removal of any edge of H leaves a path. Given a cycle H , we may order its vertices according to some arbitrarily chosen traversal of its edges; call an edge of the cycle *positive* or *negative* according to this

choice. A cycle is *balanced* if it has the same number of positive and negative edges, and *unbalanced* otherwise.

Before we state and prove our result, we need an auxiliary lemma about *multi-sorted* CSP's; here we follow Bulatov and Jeavons' treatment closely (see [9] and [10]). Let A_1, \dots, A_n be non-empty finite sets. Let $1 \leq i_1, \dots, i_k \leq n$; if $\theta \subseteq A_{i_1} \times \dots \times A_{i_k}$, we will say that θ is a *k-ary relation over* $\{A_1, \dots, A_n\}$ and has *signature* (i_1, \dots, i_k) . Let $\Gamma = \{\theta_1, \dots, \theta_s\}$ be a collection of relations over $\{A_1, \dots, A_n\}$. Then define the following *multi-sorted constraint satisfaction problem*:

- $\text{CSP}(\Gamma)$

Input: A structure $\langle I; \mu_1, \dots, \mu_s \rangle$ where each μ_i has the same arity as θ_i .

Question: Is there a map $\phi: I \rightarrow \cup A_i$ such that $\phi(\mu_i) \subseteq \theta_i$ for all i ?

Bulatov and Jeavons associate an algebra to each multi-sorted CSP as follows: let $A = A_1 \times \dots \times A_n$, and for each k -ary relation θ over $\{A_1, \dots, A_n\}$ with signature (i_1, \dots, i_k) , define a k -ary relation $\chi(\theta)$ over A by

$$\chi(\theta) = \{(a_1, \dots, a_k) \in A^k : (a_1[i_1], \dots, a_k[i_k]) \in \theta\}$$

where $a[j]$ indicates the j -th coordinate of the tuple a . Let $\widehat{\Gamma} = \Gamma \cup \{=_1, \dots, =_s\}$ where for each i , $=_i$ denotes the binary equality relation on A_i . Finally, let $\chi(\widehat{\Gamma}) = \{\chi(\rho) : \rho \in \widehat{\Gamma}\}$. Then the algebra $\mathbb{A}_{\chi(\widehat{\Gamma})}$ has universe A and its basic operations are all those that preserve every relation in $\chi(\widehat{\Gamma})$.

The following is a generalisation of Theorem 10 of [15].

Lemma 5.1. *Let A_1, \dots, A_n be 2-element sets and let Γ be a set of relations over $\{A_1, \dots, A_n\}$. If the algebra $\mathbb{A}_{\chi(\widehat{\Gamma})}$ has a Taylor term, then the multi-sorted problem $\text{CSP}(\Gamma)$ is in \mathbf{P} .*

Proof. It is easy to verify that because of the relations $\chi(=_i)$, the algebra $\mathbb{A}_{\chi(\widehat{\Gamma})}$ is isomorphic to a product $\mathbb{A}_1 \times \dots \times \mathbb{A}_n$ of algebras with universes A_1, \dots, A_n respectively. Let t be a Taylor term of $\mathbb{A}_{\chi(\widehat{\Gamma})}$ and let t^i denote the corresponding Taylor term of the algebra \mathbb{A}_i . Then inspection of the Post lattice (see for instance [28] page 36) shows that the clone generated by t^i contains either a semilattice term, a near-unanimity term or an affine term $m(x, y, z) = x + y + z$. It follows from Theorems 3.12 and 4.3 of [10] that $\text{CSP}(\Gamma)$ is in \mathbf{P} . \square

Feder shows that every oriented graph H which is a path or an unbalanced cycle admits a majority operation, and hence in this case $\text{CSP}(H)$ has bounded width. Obviously if a cycle H is not a core, then it retracts onto a path, and then $\text{CSP}(H)$ has bounded width by Lemma 4.3. It thus suffices to consider the case where H is a core.

Proposition 5.2. *Let H be an irreflexive, oriented cycle which is a core, and let \mathbb{A} be an algebra for H .*

- (1) *If $\mathcal{V}(\mathbb{A})$ admits type 1, then $\text{CSP}(H)$ is **NP**-complete, and it does not have bounded width.*
- (2) *If $\mathcal{V}(\mathbb{A})$ omits type 1, then $\text{CSP}(H)$ is in **P**, and in fact has bounded width.*

Proof. (1) It follows from the result of [11] mentioned just before Lemma 4.1 that if $\mathcal{V}(\mathbb{A})$ admits type 1, then $\text{CSP}(H)$ is **NP**-complete, and it follows from Theorem 4.2 and Lemma 4.3 that it does not have bounded width.

(2) Suppose that $\mathcal{V}(\mathbb{A})$ omits type 1. Let \mathbb{B} denote the idempotent reduct of \mathbb{A} . We may assume by the preceding remarks that H is a balanced cycle. It is then possible to partition the vertex set of H into levels, numbered $0, 1, \dots, r$, in such a way that if (x, y) is an edge of H , then the level of y is equal to one more than the level of x . Let l denote the lowest level (level 0) and let h denote the highest level (level r). Define an equivalence relation α on l as follows: vertices x, y of l satisfy $(x, y) \in \alpha$ if one may traverse the cycle H from x to y without going through level h . Define β similarly on h . Feder shows that if α and β have one block each, then H admits a majority operation, so in this case we are done. Inspection of his analysis when α and β each have at least 3 blocks shows that in this case $\mathcal{V}(\mathbb{A})$ admits type 1; indeed, all reductions used are actually constructions of subalgebras of powers of \mathbb{B} . Hence we may now suppose that α and β each have exactly 2 blocks (in [15] this case is denoted by $l^+h^+l^+h^+$).

It follows from Feder's use of "generic paths", that (i) l and h are subuniverses of \mathbb{B} and (ii) α and β are congruences of the corresponding subalgebras. Let \mathbb{B}_1 and \mathbb{B}_2 denote the 2-element quotient algebras obtained; let $B_1 = \{0_l, 1_l\}$ denote the universe of \mathbb{B}_1 and let $B_2 = \{0_h, 1_h\}$ denote the universe of \mathbb{B}_2 . Now Feder's proof shows that there exists a set $\Gamma = \{\theta_1, \theta_2, \dots, \theta_s\}$ of relations over $\{B_1, B_2\}$, where each θ_i is a subuniverse of $\mathbb{B}_1 \times \mathbb{B}_2$, $\mathbb{B}_1 \times \mathbb{B}_1 \times \mathbb{B}_2$ or $\mathbb{B}_1 \times \mathbb{B}_2 \times \mathbb{B}_2$, and such that $\text{CSP}(H)$ reduces in polynomial time to the multi-sorted $\text{CSP}(\Gamma)$.

By Lemma 9.4 and Theorem 9.6 of [19], if the variety generated by \mathbb{B} omits type 1, then the algebra \mathbb{B} has a Taylor term. Thus $\mathbb{B}_1 \times \mathbb{B}_2$ has a Taylor term, and a simple computation shows that it preserves $\chi(\theta_i)$ for each i ; since by definition it preserves $\chi(=1)$ and $\chi(=2)$, it is a term of the algebra $\mathbb{A}_{\chi(\hat{\Gamma})}$. It follows by the last lemma that $\text{CSP}(\Gamma)$ is in **P**, and hence $\text{CSP}(H)$ is also in **P**. Finally, in [15] Feder shows that all tractable cases of $\text{CSP}(H)$ have bounded width, and this completes the proof. \square

5.2. Finite order-primal algebras. For our second application we need to describe in detail a construction in [17], where Feder and Vardi associate to every

relational structure B a poset P of depth 3, to prove that every CSP is polynomial-time equivalent to a poset retraction problem. The construction is in 3 steps: first a bipartite graph H is constructed from B ; then H is modified to obtain a graph H' which is domination-free; finally the poset P is defined from H' .

It is easy to see that the homomorphism problem for any structure B of finite type is equivalent to a problem for a structure on the same base set but with a single relation: simply take the product of the basic relations. So we assume in what follows that B has an m -element base set $B = \{b_1, \dots, b_m\}$ and a single nonempty basic relation R with arity $k \geq 1$. B is a *core* if the unary operations in the clone of B are permutations. Feder and Vardi show that if B is a core, then the problems $\text{CSP}(B)$ and $\text{Ret}(H)$ are polynomial-time equivalent. The bipartite graph H is shown in Figure 1. Its vertices consist of (i) the elements b_1, \dots, b_m of B , (ii) the elements r_1, \dots, r_l of R , (iii) for each $i = 1, \dots, k$ an additional copy (of the base set) of B , depicted by black patches in the figure, and (iv) extra vertices $c_1, \dots, c_k, b_0, B', R_0$ and R' . The adjacency relation in H is defined as follows: for each $i = 1, \dots, k$, the vertex c_i is adjacent to all of the elements of the i -th copy of B , and b_0 is adjacent to all of the elements in these copies of B . The copies of the j -th element of B are adjacent to the vertex b_j , $j = 1, \dots, m$. All of the vertices r_1, \dots, r_l are adjacent to the vertices R_0 and R' . Each r_s is adjacent to the i -th copy of its i -th component, $i = 1, \dots, k$ (i.e., if $r_s = (d_1, \dots, d_k)$, then r_s is adjacent to the copy of d_i in the i -th copy of B). Each b_j is adjacent to B' , $j = 1, \dots, m$. Finally, each b_j , $j = 0, \dots, m$, is adjacent to R_0 . The colour classes of H are denoted by S and T . Our assumption on B implies that H is connected, and both of S and T have at least 3 elements.




FIGURE 1. The bipartite graph H .

In a graph, a vertex u *dominates* a vertex v if every neighbour of v is also a neighbour of u . Starting from H Feder and Vardi define a new domination-free

bipartite graph H' such that $\text{Ret}(H)$ and $\text{Ret}(H')$ are polynomial-time equivalent. Let the elements of the colour classes S and T of H be denoted by s_1, \dots, s_k and t_1, \dots, t_l , respectively. The graph H' is the incidence graph of H supplemented by additional elements r, s, s', t, t' and two copies of the 12-element graph H_1 shown in Figure 2, glued to s and t at $1'$. The adjacency relation of H' is shown in Figure 3. The vertices e_1, \dots, e_m represent each an edge of H and e_i is adjacent to its H -endpoints and r in H' ; s and s' are adjacent to s_1, \dots, s_k and s' is also adjacent to r , and similarly t and t' are adjacent to t_1, \dots, t_l , and t' is also adjacent to r .




FIGURE 2. The bipartite graph H_1 .




FIGURE 3. The bipartite graph H' .




FIGURE 4. The Feder-Vardi poset P .

In [17], Feder and Vardi prove that every bipartite graph retraction problem is polynomial-time equivalent to a poset retraction problem. Moreover, at the end of their proof they show that the poset may be chosen to be of depth 3. In the following definition the poset of depth 3 constructed for a bipartite graph is taken from the proof of Theorem 15 of [17].

The *Feder–Vardi poset* P for the relational structure B is the poset obtained as follows from H' : the base set of P consists of the maximal complete bipartite induced subgraphs M of H' and the ordering of P is defined by containment of the sets $M \cap T'$, where T' is the colour class of H' containing t (the lower level, in the figure.) A straightforward computation yields the poset depicted in Figure 4. Simply note that s_i^* (t_j^*) corresponds to the maximal complete bipartite induced subgraph M with $M \cap S' = \{s_i, r\}$ ($\{t_j, r\}$) where S' is the colour class of H' containing r (the upper level in the figure.)

Observe that the subgraph spanned by $s_1^*, \dots, s_k^*, t_1^*, \dots, t_l^*$, and e_1, \dots, e_m in Figure 4 is an isomorphic copy of the incidence graph of H . Moreover, P_1 and P'_1 in Figure 4 are two copies of the poset obtained by applying to H_1 the construction just described for H' .

Theorem 5.3. *Let B be a relational structure with a single relation and \mathbb{B} an algebra for B . Let \mathcal{P} be the relational structure obtained from the Feder–Vardi poset P related to B by adding all one-element subsets of P as unary relations. Let \mathbb{A} be an algebra for \mathcal{P} . Then $\mathcal{V}(\mathbb{A})$ interprets in $\mathcal{V}(\mathbb{B})$.*

Proof. We use the notation introduced above for the construction of the graphs H , H' and H_1 . Clearly the basic (or equivalently, term) operations of the algebra \mathbb{A} are the idempotent, order-preserving (monotone) operations on P . Similarly, the operations of \mathbb{B} are the structure-preserving operations on B . We define a map α that assigns to every idempotent monotone operation f on P a term operation of \mathbb{B} . Since $\alpha(f)$ will be defined as a restriction of f to an appropriate subset of P , it will follow immediately that the map α is a clone homomorphism, and this will conclude the proof of the theorem. We define α in two steps that follow the construction of the graphs H and H' defined earlier. We construct a partial map f' on H and then a map $\alpha(f) = f''$ on B in such a way that both maps are idempotent, structure preserving and obtained by restriction.

Let f be an n -ary monotone idempotent operation on P . (1) We define f' and show that it is structure preserving and idempotent. First we require some subalgebras of \mathbb{A} , which we construct using primitive positive formulas of A . Let $E = \{e_1, \dots, e_m\}$, $S^* = \{s_1^*, \dots, s_k^*\}$ and $T^* = \{t_1^*, \dots, t_l^*\}$. Then

$$X_S = \{x : x \leq r \text{ and } \exists z \text{ such that } x \leq z \geq s\} = S^* \cup E \cup \{s'\}$$

and similarly $X_T = T^* \cup E \cup \{t\}$ are subalgebras of \mathbb{A} . Hence so is their intersection E . Moreover,

$$Y_S = \{x \in X_S : \exists u \in E \text{ such that } u \leq x\} = S^* \cup E$$

and similarly $Y_T = T^* \cup E$ are also subalgebras of \mathbb{A} .

If u and v are adjacent in H , where $u \in S$ and $v \in T$, let e_{uv} denote the edge determined by u and v in E . We define two binary relations on E as follows:

$$\begin{aligned} e \sim_S e', & \text{ if there exists } z \in Y_S \text{ such that } e \leq z \leq e', \\ e \sim_T e', & \text{ if there exists } z \in Y_T \text{ such that } e \leq z \leq e'. \end{aligned}$$

Both relations are subalgebras of E^2 since they are defined by primitive positive formulas, and in fact they are congruences of E . Indeed, $e_{uv} \sim_S e_{u'v'}$ if and only if $u = u'$ and $e_{uv} \sim_T e_{u'v'}$ if and only if $v = v'$. It follows that f induces an operation on E such that

$$f(e_{u_1v_1}, \dots, e_{u_nv_n}) = e_{\phi(u_1, \dots, u_n)\psi(v_1, \dots, v_n)},$$

where $\phi: S^n \rightarrow S$ and $\psi: T^n \rightarrow T$ are appropriate maps. Since f is idempotent, $f(e_{uv}, \dots, e_{uv}) = e_{uv}$ and hence ϕ and ψ are also idempotent.

We show that $f|_S = \phi$; the proof that $f|_T = \psi$ is similar. Observe that

$$S = \{x \in P : x \geq s^* \text{ and } \exists u \in E \text{ such that } u \leq x\},$$

and hence S is a subalgebra of \mathbb{A} . Let $s_{i_m} \in S$ and $t_{j_m} \in T$ be adjacent for $m = 1, \dots, n$. Then

$$f(s_{i_1}, \dots, s_{i_n}) \geq f(s_{i_1}^*, \dots, s_{i_n}^*) \geq f(e_{s_{i_1}t_{j_1}}, \dots, e_{s_{i_n}t_{j_n}}) = e_{uv}$$

for some $u \in S$ and $v \in T$. But the only element of S which is above e_{uv} is u itself, hence $f(s_{i_1}, \dots, s_{i_n}) = u = \phi(s_{i_1}, \dots, s_{i_n})$.

Let $f' = f|_{S^n \cup T^n}$. Then $f'(S^n) \subseteq S$ and $f'(T^n) \subseteq T$ since S and T are subalgebras of \mathbb{A} . Furthermore, for any $s_{i_m} \in S$ and $t_{j_m} \in T$ that are adjacent, $m = 1, \dots, n$, we have that

$$f'(s_{i_1}, \dots, s_{i_n}) = f(s_{i_1}, \dots, s_{i_n}) = \phi(s_{i_1}, \dots, s_{i_n})$$

is adjacent to

$$f'(t_{j_1}, \dots, t_{j_n}) = f(t_{j_1}, \dots, t_{j_n}) = \psi(t_{j_1}, \dots, t_{j_n})$$

by definition of ϕ and ψ . Thus, f' is an edge preserving and idempotent partial map on H .

(2) To complete the proof we define the operation f'' on B and show that it is structure preserving. Notice that in H , the set of neighbours of B' is precisely $B = \{b_1, \dots, b_m\}$, and since f' is idempotent and edge-preserving it must preserve this set. We define f'' to be the restriction of f' to B^n .

In order to prove that f'' preserves R we consider k -tuples

$$\sigma_1 = (s_{1,1}, \dots, s_{k,1}), \dots, \sigma_n = (s_{1,n}, \dots, s_{k,n})$$

in R . For $j = 1, \dots, k$, let C_j denote the neighborhood of c_j in H and for any $u \in B$ let u^j denote its copy in C_j . By using the idempotency of f' again we obtain

$$f'(s_{j,1}^j, \dots, s_{j,n}^j) \in C_j \text{ for any } j = 1, \dots, k \text{ and } f'(\sigma_1, \dots, \sigma_n) \in R \text{ in } H.$$

On the other hand, $f'(\sigma_1, \dots, \sigma_n)$ and $f'(s_{j,1}^j, \dots, s_{j,n}^j)$ are adjacent for any $j = 1, \dots, k$. Hence $[f'(\sigma_1, \dots, \sigma_n)]^j = f'(s_{j,1}^j, \dots, s_{j,n}^j)$ for all $j = 1, \dots, k$. Moreover, $f'(s_{j,1}^j, \dots, s_{j,n}^j)$ and $f'(s_{j,1}, \dots, s_{j,n})$ are adjacent for all j ; by the definition of H this means that the k -tuple $(f'(s_{1,1}, \dots, s_{1,n}), \dots, f'(s_{k,1}, \dots, s_{k,n}))$ is in R , which concludes the proof. \square

The following corollary is a straightforward consequence of Theorems 3.4 and 5.3.

Corollary 5.4. *Let B be a relational structure with a single relation, and let P be its Feder–Vardi poset. If $\text{Ret}(P)$ has bounded width, then $\text{CSP}(B)$ also has bounded width.*

We do not know if the converse of Corollary 5.4 holds. We note however that bounded strict width of structures is never preserved by the Feder–Vardi construction. The problem $\text{CSP}(B)$ has *strict width* l if, for any partial map f from a structure I to B that does not extend to a full homomorphism, there exists a subset of its domain with at most l elements such that the restriction of f to this subset still does not extend ([17] page 82, see also [4]). No matter what B is, its Feder–Vardi poset P is always a *ramified poset*, that is, P is a connected poset with at least two elements and has no element with a unique lower or upper cover. It was shown in [24] that no ramified poset admits a near unanimity operation, hence by Theorem 25 in [17] $\text{CSP}(\mathcal{P})$, and in particular $\text{Ret}(P)$, does not have bounded strict width.

Until now, there was no known example of a poset P such that the variety generated by an algebra for \mathcal{P} admits type 2 but omits type 1. By using the Feder–Vardi construction we can present such an example, although for the time being we need to assume that $\mathbf{P} \neq \mathbf{NP}$.

Proposition 5.5. *Let P be the Feder–Vardi poset of the two element structure $B = (\{0, 1\}; \{(x, y, z, 0) : x + y + z = 1\})$. The variety generated by an algebra \mathbb{A} for \mathcal{P} admits type 2 and omits type 1, provided $\mathbf{P} \neq \mathbf{NP}$.*

Proof. Clearly the clone of \mathbb{B} is the set of idempotent operations of the two element vector space. By the previous theorem $\mathcal{V}(\mathbb{A})$ interprets in $\mathcal{V}(\mathbb{B})$, where \mathbb{B} is an algebra for B . So $\mathcal{V}(\mathbb{A})$ interprets in a variety generated by the two element vector space, an affine algebra. Hence by Lemma 4.1, $\mathcal{V}(\mathbb{A})$ admits type 1 or 2. Moreover,

$\mathcal{V}(\mathbb{A})$ omits type 1, for otherwise $\mathcal{V}(\mathbb{A})$ would interpret in the variety of sets by Theorem 9.6 in [19] and so $Ret(P)$ would be **NP**-complete, which is impossible as both $CSP(B)$ and $Ret(P)$ are polynomial-time solvable by Theorem 14 and Theorem 32 in [17]. \square

6. Conclusion

We have shown in Theorem 4.2 that, if the algebra associated to a CSP of bounded width is idempotent, then it generates a variety that omits types 1 and 2; Lemma 4.3 shows that there is no loss of generality in considering idempotent algebras, and that this result can be used to show that several decision problems do not have bounded width (section 5). We conjecture that a finite idempotent algebra \mathbb{A} has bounded width if and only if $\mathcal{V}(\mathbb{A})$ omits types 1 and 2. By Theorem 9.10 in [19] the property that the variety generated by an algebra for a finite relational structure of finite type omits types 1 and 2 is decidable. Hence, if our conjecture holds, then determining whether a finite structure of finite type has bounded width is decidable.

Our conjecture is verified for finite Boolean algebras, distributive lattices and semilattices, the building blocks of varieties omitting types 1 and 2. More generally, the conjecture is confirmed for finite idempotent algebras with a near unanimity term operation or with a totally symmetric term operation of sufficiently large arity, see [17]. We also note that the conjecture is true for finite conservative algebras by a recent result of Bulatov in [6].

In [5], Bulatov proves a result which is similar to our Theorem 4.2 and states a conjecture which is parallel to ours, but uses a slightly different notion of width. We now clarify briefly the connection. For this, it will be convenient to alter slightly our definition of CSP, but the reader will easily verify that the two approaches coincide when the set Γ of constraint relations is finite (see [5], [21]).

Let Γ be a (possibly infinite) set of finitary relations on the finite set A . Define $CSP(\Gamma)$ as the following decision problem:

Input: A pair $P = (V, C)$ where V is a finite non-empty set of *variables*, and C is a finite set of *constraints*, i.e., $C = \{(s_1, \theta_1), \dots, (s_r, \theta_r)\}$ where for each $i = 1, \dots, r$, s_i is a tuple of (not necessarily distinct) variables and θ_i is a relation in Γ of the same arity as s_i . s_i is called the *constraint scope*, and θ_i is called the *constraint relation*.

Question: Is there a *solution to P* , i.e., a map f from V to A such that, for each $i = 1, \dots, r$, if $s_i = (v_{j_1}, \dots, v_{j_t})$, then we have $(f(v_{j_1}), \dots, f(v_{j_t})) \in \theta_i$.

Let θ be a relation of arity m on A , and let t be an m -tuple of elements of A . For $1 \leq i \leq m$ we denote the i -th coordinate of t by $t[i]$. Let (i_1, \dots, i_k) be a list of (not necessarily distinct) integers between 1 and m . The *projection* $\pi_{i_1, \dots, i_k}(t)$ is defined to be the k -tuple

$$\pi_{i_1, \dots, i_k}(t) = (t[i_1], \dots, t[i_k]).$$

The *projection* $\pi_{i_1, \dots, i_k}(\theta)$ is defined to be the k -ary relation

$$\pi_{i_1, \dots, i_k}(\theta) = \{\pi_{i_1, \dots, i_k}(t) : t \in \theta\}.$$

Let $P = (V, C)$ be an instance of $\text{CSP}(\Gamma)$. Let V' be a subset of V . The *subproblem* $P|_{V'}$ is the problem instance of $\text{CSP}(\Gamma)$ with set of variables V' and the following constraints: for each constraint (s, θ) of P such that s contains some entries in V' , choose $I = (i_1, \dots, i_r)$ to be a list of the indices of the entries of s that are in V' , and make $(\pi_I(s), \pi_I(\theta))$ a constraint of $P|_{V'}$. We call a solution of the problem $P|_{V'}$ a *partial solution of P on V'* .

Let $k \geq 1$. An instance $P = (V, C)$ of $\text{CSP}(\Gamma)$ is *k -minimal* if it satisfies the following conditions: (i) every k -subset of V is contained in the scope of some constraint in C and (ii) for every constraint $(s, \theta) \in C$ and every k -subset W of V , the restriction of any tuple of θ to $s \cap W$ extends to a partial solution on W .

Let \mathbb{A} be a finite algebra and let Γ denote the set of all subuniverses of finite powers of \mathbb{A} . We say that \mathbb{A} has *relational width k* if, whenever the constraint relations of a k -minimal instance P for $\text{CSP}(\Gamma)$ are non-empty, then P has a solution. The algebra has *bounded relational width* if it has relational width k for some $k \geq 1$.

Proposition 6.1. *If a finite algebra \mathbb{A} has bounded relational width, then it has bounded width.*

Proof. Suppose that \mathbb{A} has bounded relational width. Let Γ denote the set of subuniverses of finite powers of \mathbb{A} , and let \mathcal{A} be a structure whose finitely many basic relations are elements of Γ ; let \mathcal{I} be an instance of $\text{CSP}(\mathcal{A})$. We may assume that \mathbb{A} has relational width k where k is at least as large as the maximum arity of the basic relations of \mathcal{A} , since bounded relational width implies relational width k for all k large enough. Suppose that the output relations ρ_K are non-empty when we run the (l, k) -algorithm on \mathcal{I} with $l = k - 1$. Construct an instance \mathcal{J} of $\text{CSP}(\Gamma)$ as follows: its base set is the base set I of \mathcal{I} and its constraints are the pairs (H, ρ_H) for every k -subset H of I . We claim that this instance is k -minimal. Indeed, let (H, ρ_H) be a constraint, let $\bar{a} \in \rho_H$ and let K be any k -subset of I ; we must show that the restriction of \bar{a} to $H \cap K$ extends to a partial solution on K . Because we are considering output relations of the (l, k) -algorithm, we have that (i) the restriction of \bar{a} to $H \cap K$ extends to a tuple $\bar{b} \in \rho_K$ and (ii) every such tuple is a partial solution on K . Since \mathbb{A} has relational width k there exists a solution

$f: I \rightarrow A$ to the instance \mathcal{J} . But since k is at least as large as the arity of every basic relation of \mathcal{A} , it follows that f is also a homomorphism from \mathcal{I} to \mathcal{A} . Thus if the (l, k) -algorithm yields non-empty output relations on an instance there is a solution to this instance, and so \mathbb{A} has bounded width. \square

It is open whether these conditions are actually equivalent, and this question is related to the relationship between local and global tractability (see for instance [11]). Finally we remark that by Lemma 1 of [5] and our Lemma 4.1, if an idempotent algebra \mathbb{A} generates a variety that omits types 1 and 2, then the graph $G(\Gamma)$ has no blue edges (as defined in [5]), where Γ is the set of subuniverses of finite powers of \mathbb{A} (we do not know if the converse holds.) This implies in particular that Theorem 1 of [5] follows from our Theorem 4.2.

Acknowledgments.

The authors wish to thank Matt Valeriote and Andrei Bulatov for fruitful discussions, and the anonymous referee for valuable comments.

REFERENCES

- [1] A. Atserias, *On Digraph Coloring Problems and Treewidth Duality*, in: Proceedings of the Twentieth Annual IEEE Symp. on Logic in Computer Science, LICS 2005, 2005, pp. 106–115.
- [2] G. Birkhoff, *On the structure of abstract algebras*, Proc. Camb. Philos. Soc. **31** (1935), 433–454.
- [3] V. G. Bodnarčuk, L. A. Kaluzhnin, V. N. Kotov and B. A. Romov, *Galois theory for Post algebras I-II* (russian), Kibernetika, **3** (1969), 1–10 and **5** (1969), 1–9. English version: *Cybernetics*, (1969), 243–252 and 531–539.
- [4] R. C. Brewster, T. Feder, P. Hell, J. Huang and G. MacGillivray, *Near unanimity functions and homomorphisms to reflexive graphs*, preprint, 2003.
- [5] A. Bulatov, *A Graph of a Relational Structure and Constraint Satisfaction Problems*, LICS 2004, pp. 448–457.
- [6] A. Bulatov, *Tractable conservative constraint satisfaction problems*, LICS 2003, p. 321.
- [7] A. Bulatov, *A dichotomy theorem for constraints on a three-element set*, in: Proceedings of 43rd IEEE Symposium on Foundations of Computer Science (FOCS'02), Vancouver, Canada, 2002, pp. 649–658.
- [8] A. Bulatov, *H-coloring dichotomy revisited*, *Theor. Comput. Sci.* **349** (2005), 31–39.
- [9] A. Bulatov and P. Jeavons, *Algebraic structures in combinatorial problems*, Technical report MATH-AL-4-2001, Technische Universität Dresden, 2001.
- [10] A. Bulatov and P. Jeavons, *An algebraic approach to multi-sorted constraints*, in: Proceedings of 9th International Conference on Principles and Practice of Constraint Programming (CP'03), Kinsale, Ireland, 2003.
- [11] A. Bulatov, A. Krokhin and P. Jeavons, *Constraint satisfaction problems and finite algebras*, in: Proceedings of 27th International Colloquium on Automata, Languages and Programming (ICALP'00), Geneva, Switzerland, Lecture Notes in Computer Science **1853** (2000), 272–282.
- [12] M. Clasen and M. Valeriote, *Tame Congruence Theory*, in: Lectures on algebraic model theory, B. Hart, M. Valeriote, eds. Providence, R. I., American Mathematical Society (Fields Institute Monographs) **15** (2002).

- [13] V. Dalmau and J. Pearson, *Closure functions and width 1 problems*, in: Proceedings of the 5th International Conference on Principles and Practice of Constraint Programming 1999, Lecture Notes In Computer Science **1713**, pp. 159–173.
- [14] R. Dechter, *From Local to Global Consistency*, Artificial Intelligence **55** (1992), 87–107.
- [15] T. Feder, *Classification of Homomorphisms to Oriented Cycles and of k -Partite Satisfiability*, SIAM J. Discrete Math. **14** (2001), 471–480.
- [16] T. Feder and M. Vardi, *Monotone monadic SNP and constraint satisfaction*, in: Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, 1993, pp. 612–622.
- [17] 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 of Computing **28** (1998), 57–104.
- [18] P. Hell and J. Nešetřil, *On the complexity of H -colouring*, J. Combin. Theory Ser. B **48** (1990), 92–110.
- [19] D. Hobby and R. McKenzie, *The structure of finite algebras*, Contemporary Mathematics **76**, American Mathematical Society, Providence, RI, 1988.
- [20] P. Jeavons, *On the algebraic structure of combinatorial problems*, Theoret. Comput. Sci. **200** (1998), 185–204.
- [21] P. G. Jeavons, D. A. Cohen and M. Cooper, *Constraints consistency and closure*, Artificial Intelligence, **101** (1998) 251–265.
- [22] P. G. Jeavons, D. A. Cohen and M. Gyssens, *Closure properties of constraints* Journal of the ACM **44** (1997), 527–548.
- [23] P. G. Kolaitis and M. Y. Vardi, *Conjunctive-query containment and constraint satisfaction*, in: Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems, Seattle, Washington, 1998, pp. 205–213.
- [24] B. Larose and L. Zádori, *Algebraic properties and dismantlability of finite posets*, Discrete Math. **163** (1997), 89–99.
- [25] B. Larose and L. Zádori, *The complexity of the extendibility problem for finite posets*, SIAM J. Discrete Math. **17** (2003), 114–121.
- [26] J. Nešetřil and X. Zhu, *On bounded treewidth duality of graphs*, Journal of Graph Theory **23**, 151–162, 1996.
- [27] T. J. Schaefer, *The complexity of satisfiability problems*, in: Proceedings of the 10th annual ACM symposium on Theory of computing (STOC), San Diego, California, 1978, pp. 216–226.
- [28] Á. Szendrei, *Clones in universal algebra*, Sém. de Mathématiques Supérieures **99**, Séminaire Scientifique OTAN, les presses de l'Université de Montréal, 1986.
- [29] Á. Szendrei, *Idempotent algebras with restrictions on subalgebras*, Acta Sci. Math. **51** (1987), 251–268.

BENOIT LAROSE

Department of Mathematics and Statistics, Concordia University, 1455 de Maisonneuve
West, Montréal, Qc, Canada, H3G 1M8
e-mail: larose@mathstat.concordia.ca
URL: <http://cicma.mathstat.concordia.ca/faculty/larose/>

LÁSZLÓ ZÁDORI

Bolyai Intézet, Aradi vértanúk tere 1, H-6720, Szeged, Hungary
e-mail: zadori@math.u-szeged.hu