

Dualities for Constraint Satisfaction Problems

Andrei Bulatov¹, Andrei Krokhin², and Benoit Larose³

¹ School of Computing Science
Simon Fraser University
Burnaby, BC, Canada, V5A 1S6
`abulatov@cs.sfu.ca`

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

³ Department of Mathematics and Statistics
Concordia University
Montréal, Qc, Canada, H3G 1M8
`larose@mathstat.concordia.ca`

Abstract. In a nutshell, a duality for a constraint satisfaction problem equates the existence of one homomorphism to the non-existence of other homomorphisms. In this survey paper, we give an overview of logical, combinatorial, and algebraic aspects of the following forms of duality for constraint satisfaction problems: finite duality, bounded pathwidth duality, and bounded treewidth duality.

1 Introduction

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, computer science, discrete mathematics, and elsewhere [18, 32, 58]. An instance of the constraint satisfaction problem is represented by a finite set V of variables, a (finite) domain D of values for each variable, and a set of constraints $\{(\bar{s}_1, R_1), \dots, (\bar{s}_q, R_q)\}$. Each constraint consists of a constraint scope \bar{s}_i , which is an m_i -tuple of variables, and a constraint relation $R_i \subseteq D^{m_i}$. The aim is then to decide whether there is an assignment $h : V \rightarrow D$ that satisfies the constraints, i.e., such that $h(\bar{s}_i) \in R_i$ for all i .

It has been observed [26] (see also [40]) 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} ? One of the most studied restrictions on the CSP is the non-uniform CSP – 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 various versions of k -SAT, GRAPH COLOURING, and SYSTEMS OF EQUATIONS (see [16, 32, 40, 48]). Strong motivation for studying this framework was given in [26] where it was shown that such problems can be used in attempts to identify a largest natural subclass of \mathbf{NP} that avoids problems of intermediate complexity.

The two main general classification problems about the class of problems of the form $\text{CSP}(\mathbf{B})$ are:

1. classify the problems $\text{CSP}(\mathbf{B})$ with respect to computational complexity, that is, for a given complexity class \mathcal{K} , characterise (under suitable complexity-theoretic assumptions) structures \mathbf{B} such that $\text{CSP}(\mathbf{B})$ is in \mathcal{K} ;
2. classify the problems $\text{CSP}(\mathbf{B})$ with respect to descriptive complexity, that is, for a given logic L , characterise structures \mathbf{B} such that $\text{CSP}(\mathbf{B})$, as the class of all structures admitting a homomorphism to \mathbf{B} , is definable in L .

In addition, there is a so-called meta-problem:

3. Determine the (computational) complexity of deciding whether, for a given structure \mathbf{B} , $\text{CSP}(\mathbf{B})$ has a certain (computational or descriptive) 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., [46]), combinatorics (e.g., [30, 32, 49]), universal algebra (e.g., [7, 10, 12, 16, 39, 48]), or combinations of those (e.g., [2, 8, 19, 26, 50]). In this survey, we will discuss a combinatorial idea that has a bearing on all the above problems, and has strong links with the three approaches — the idea of *homomorphism duality*.

The concept of duality has been much used to study homomorphism problems. In essence, a duality equates the existence of one homomorphism to the non-existence of some other homomorphism(s). 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 $\mathcal{O}_{\mathbf{B}}$. Of course, the set $\mathcal{O}_{\mathbf{B}}$ can always be chosen to consist of all structures that do not homomorphically map to \mathbf{B} , but this choice does not give any information about $\text{CSP}(\mathbf{B})$. If, however, $\mathcal{O}_{\mathbf{B}}$ can be chosen so that it has certain nice properties then this can tell us much about the computational or descriptive complexity of $\text{CSP}(\mathbf{B})$.

Most of the early studies of dualities were restricted to the case of (di)graphs (see survey [34], also [32, 33, 36, 37, 47, 55]). For general relational structures, the main forms of duality that have been considered in the literature are finite duality, bounded pathwidth duality, and bounded treewidth duality. We give the necessary combinatorial, logical, and algebraic preliminaries in Section 2, and then consider the three dualities in Sections 3, 4, and 5, respectively. Sections 6 and 7 contain some remarks and a list of open problems concerning dualities.

2 Preliminaries

2.1 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 . Let $\maxar(\mathbf{A})$ denote the maximum arity of a relation in \mathbf{A} . Unless specified otherwise, 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 $h : \mathbf{A} \rightarrow \mathbf{B}$, and the set of all homomorphisms from \mathbf{A} to \mathbf{B} is denoted by $\text{hom}(\mathbf{A}, \mathbf{B})$. We also say that \mathbf{A} homomorphically maps to \mathbf{B} , and write $\mathbf{A} \rightarrow \mathbf{B}$ if there is a homomorphism from \mathbf{A} to \mathbf{B} and $\mathbf{A} \not\rightarrow \mathbf{B}$ if there is no homomorphism. Now $\text{CSP}(\mathbf{B})$ can be defined to be the class of all structures \mathbf{A} such that $\mathbf{A} \rightarrow \mathbf{B}$. The class of all structures \mathbf{A} such that $\mathbf{A} \not\rightarrow \mathbf{B}$ will be denoted by $\text{co-CSP}(\mathbf{B})$.

Example 1. If \mathbf{B}_{hc} is a digraph \mathbf{H} then $\text{CSP}(\mathbf{B}_{hc})$ is the much-studied problem, **H-COLOURING**, of deciding whether there is a homomorphism from a given digraph to \mathbf{H} [32]. If \mathbf{H} is the complete graph \mathbf{K}_k on k vertices then it is well known (and easy to see) that $\text{CSP}(\mathbf{B}_{hc})$ is precisely the k -COLOURING problem.

Example 2. If \mathbf{B}_{lhc} is a structure obtained from a digraph \mathbf{H} by adding, for each non-empty subset U of H , a unary relation U then $\text{CSP}(\mathbf{B}_{lhc})$ is exactly the **LIST H-COLOURING** problem, in which every vertex v of the input digraph \mathbf{G} gets a list L_v of vertices of \mathbf{H} , and the question is whether there is a homomorphism $h : \mathbf{G} \rightarrow \mathbf{H}$ such that $h(v) \in L_v$ for all $v \in G$ (see [32]).

Example 3. If \mathbf{B}_{unr} is the Boolean (i.e., with universe $\{0, 1\}$) structure with one binary relation Eq , which is the equality relation, and two unary relations $\{0\}$ and $\{1\}$ then $\text{CSP}(\mathbf{B}_{unr})$ is the (undirected) **UNREACHABILITY** problem where one is given a graph and two sets of vertices in it, S and T , and the question is whether there is no path in the graph from any vertex in S to a vertex in T .

Example 4. In the **PATH SYSTEM ACCESSIBILITY** problem [29], one is given a relational structure \mathbf{A} with one ternary relation $P^{\mathbf{A}}$, and two unary relations $S^{\mathbf{A}}$ and $T^{\mathbf{A}}$. The unary relations represent “source” and “terminal” nodes, respectively. The question is whether there is an “accessible” terminal node, where a node x is accessible if $x \in S^{\mathbf{A}}$ or $(a, b, x) \in P^{\mathbf{A}}$ for some accessible $a, b \in A$.

Let \mathbf{B}_{ps} be the Boolean structure with one ternary relation $P^{\mathbf{B}_{ps}} = \{(x, y, z) \mid x \wedge y \rightarrow z\}$ and two unary relations $S^{\mathbf{B}_{ps}} = \{1\}$ and $T^{\mathbf{B}_{ps}} = \{0\}$. Then it is easy to verify that the **PATH SYSTEM ACCESSIBILITY** problem is precisely $\text{co-CSP}(\mathbf{B}_{ps})$.

Example 5. Let \mathbf{B}_{3H} be the structure with universe $\{0, 1\}$, one unary relation $U^{\mathbf{B}_{3H}} = \{1\}$ and two ternary relations $P^{\mathbf{B}_{3H}} = \{0, 1\}^3 \setminus \{(1, 1, 0)\}$ and $N^{\mathbf{B}_{3H}} = \{0, 1\}^3 \setminus \{(1, 1, 1)\}$. It is easy to see that every Horn 3-CNF formula φ with variables x_1, \dots, x_n can be represented as a structure \mathbf{A}_φ with universe $\{x_1, \dots, x_n\}$ and relations $U^{\mathbf{A}_\varphi}$, $P^{\mathbf{A}_\varphi}$, $N^{\mathbf{A}_\varphi}$ where $U^{\mathbf{A}_\varphi}$ is the set of all unit clauses (in φ), $P^{\mathbf{A}_\varphi}$ is the set of all clauses of the form $(\neg x \vee \neg y \vee z)$, and $N^{\mathbf{A}_\varphi}$ is the set of all clauses of the form $(\neg x \vee \neg y \vee \neg z)$. Clearly, we have $\mathbf{A}_\varphi \rightarrow \mathbf{B}_{3H}$ if and only if φ is satisfiable. Hence HORN 3-SAT is precisely $\text{CSP}(\mathbf{B}_{3H})$.

Example 6. Let \mathbf{B}_{le} be a structure with universe $\{0, 1\}$, one ternary relation $\{(x, y, z) \mid x + y + z = 1 \pmod{2}\}$, and one unary relation $\{0\}$. It is well known, and easy to verify, that $\text{CSP}(\mathbf{B}_{le})$ is the problem of solving systems of linear equations (with at most 3 variables per equation) over the two-element field.

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} . A *projective* homomorphism from \mathbf{A} to \mathbf{B} is a partial mapping h from A to B such that, for any $R \in \tau$ (say, of arity n) and any tuple $(a_1, \dots, a_n) \in R^{\mathbf{A}}$, there exists a tuple $(b_1, \dots, b_n) \in R^{\mathbf{B}}$ such that $h(a_i) = b_i$ for every a_i in $\text{dom}(h)$, the domain of h . Clearly, every projective homomorphism is also a partial homomorphism.

A *retract* of a structure \mathbf{B} is an induced substructure \mathbf{B}' of \mathbf{B} such that there is a homomorphism $g : \mathbf{B} \rightarrow \mathbf{B}'$ with $g(b) = b$ for every $b \in B'$. In this case we (trivially) have that $\text{CSP}(\mathbf{B})$ and $\text{CSP}(\mathbf{B}')$ coincide. A structure is called a *core* if it has no homomorphism to any of its proper substructures. A retract of \mathbf{B} that has minimal size among all retracts of \mathbf{B} is called a *core* of \mathbf{B} . It is well known that all cores of a structure are isomorphic, and so one speaks of *the core*, $\text{core}(\mathbf{B})$, of a structure \mathbf{B} .

2.2 Obstructions and dualities

In order to define some of our dualities, we will need the notions of pathwidth and treewidth of relational structures.

Definition 1. For $0 \leq j \leq k$, 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) .

Example 7. 1. Consider the graph \mathbf{G} from Fig. 1. The top-left decomposition shows that \mathbf{G} has treewidth at most $(1, 3)$, the top-right and the bottom decompositions imply that \mathbf{G} has pathwidth at most $(1, 5)$ and at most $(2, 4)$, respectively.




Fig. 1. Examples of pathwidth and treewidth.

2. Any cycle has pathwidth at most (2,3). Indeed, assume that the nodes of the cycle are $0, 1, \dots, n-1$ and the edges are $(i, i+1)$, where addition is modulo n . Consider a path with nodes S_1, S_2, \dots, S_{n-2} where $S_i = \{0, i, i+1\}$. It is easy to check that this is a path-decomposition of the cycle.
3. Any tree has treewidth at most (1,2). Indeed, take T to have the edges of the original tree as nodes and the adjacency relation given by the incidence relation of edges in the original tree.

Note that we use two numbers to parameterise treewidth and pathwidth, as is customary in the study of CSPs [19, 26, 52] (rather than one as is customary in graph theory [32]), for the following reason. The first parameter j gives a more convenient parameterisation of CSPs, since the second parameter k 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 vocabularies 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 from graph theory, as follows. The *Gaifman graph* $\mathbf{G}(\mathbf{A})$ of a structure \mathbf{A} is defined to have the same universe (set of vertices) as \mathbf{A} and the edges of $\mathbf{G}(\mathbf{A})$ are the pairs (a, a') of distinct elements such that a and a' appear in the same tuple in some relation in \mathbf{A} . Then it is not hard to check that the following numbers are equal:

- the minimum k such that \mathbf{A} has pathwidth (treewidth) at most $(k, k+1)$,
- pathwidth (treewidth, respectively) of $\mathbf{G}(\mathbf{A})$ in the sense of graph theory.

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 $\mathbf{A}' \in \mathcal{O}$.

Note that sometimes such sets are called “complete obstruction sets”.

Definition 3. A structure \mathbf{B} is said to have finite duality if it has a finite obstruction set.

Example 8. Let \mathbf{T}_n be the *transitive tournament* on n vertices, that is, the universe of \mathbf{T}_n is $\{0, 1, \dots, n-1\}$, and the only relation is the binary relation $\{(i, j) \mid 0 \leq i < j \leq n-1\}$. Also, let \mathbf{P}_n be the *directed path* on $n+1$ vertices, that is the structure with universe $\{0, 1, \dots, n\}$ and the relation $\{(i, i+1) \mid 0 \leq i \leq n-1\}$. It is well known (see, e.g., Proposition 1.20 of [32]) and easy to show that, for any digraph \mathbf{G} , $\mathbf{G} \rightarrow \mathbf{T}_n$ if and only if $\mathbf{P}_n \not\rightarrow \mathbf{G}$. Hence, $\{\mathbf{P}_n\}$ is an obstruction set for \mathbf{T}_n , and \mathbf{T}_n has finite duality.

Definition 4. A τ -structure \mathbf{B} is said to have (j, k) -pathwidth duality⁴ if it has an obstruction set consisting of structures of pathwidth at most (j, k) . In other words, \mathbf{B} has (j, k) -pathwidth duality if, for any τ -structure \mathbf{A} , we have $\mathbf{A} \rightarrow \mathbf{B}$ if and only if $\mathbf{C} \rightarrow \mathbf{A}$ implies $\mathbf{C} \rightarrow \mathbf{B}$ for every τ -structure \mathbf{C} of pathwidth at most (j, k) .

We say that \mathbf{B} has j -pathwidth duality if it has (j, k) -pathwidth duality for some $k \geq j$, and \mathbf{B} has bounded pathwidth duality if it has j -pathwidth duality for some $j \geq 0$.

Example 9. It is well known that a graph \mathbf{G} is 2-colourable if and only if it contains no odd cycles, which is the same as to say that \mathbf{G} does not admit a homomorphism from any odd cycle. Since the 2-COLOURABILITY problem is the same as $\text{CSP}(\mathbf{K}_2)$, we obtain that the family of all odd cycles forms an obstruction set for \mathbf{K}_2 . By Example 7, any cycle has pathwidth at most $(2, 3)$, so the structure \mathbf{K}_2 has $(2, 3)$ -pathwidth duality. It is easy to see that \mathbf{K}_2 does not have finite duality.

Definition 5. By replacing “pathwidth” with “treewidth” throughout Definition 4, one obtains the corresponding definitions of treewidth dualities.

Example 10. The structure \mathbf{B}_{ps} from Example 4 has $(1, 3)$ -treewidth duality. To prove this, we need to show that, for any structure $\mathbf{A} \in \text{co-CSP}(\mathbf{B}_{ps})$, there exists a structure $\mathbf{C} \in \text{co-CSP}(\mathbf{B}_{ps})$ such that $\mathbf{C} \rightarrow \mathbf{A}$ and \mathbf{C} has treewidth at most $(1, 3)$. If $\mathbf{A} \in \text{co-CSP}(\mathbf{B}_{ps})$ then we can choose some terminal node in \mathbf{A} that can be “accessed” (or “derived”) from the source nodes. It is clear that this derivation procedure can be represented as a “tree”, as shown in Fig. 2. The substructure \mathbf{A}' of \mathbf{A} (corresponding to the derivation) is shown on the right; d and e are source nodes, t is a terminal node, and every oval depicts a unit derivation via a triple from the relation $P^{\mathbf{A}}$. Now modify the structure \mathbf{A}' as follows: for every element $x \in A'$, give new names to the occurrences of x in A' so that each element in the obtained structure appears either in a single oval or else in two ovals such that this element is the intersection of the two ovals, and then modify the set of source nodes accordingly. Let \mathbf{C} be the obtained structure

⁴ Called (j, k) -path duality in [19].

(see Fig. 2, left). It is clear that \mathbf{C} has treewidth at most (1,3). Furthermore, we have $\mathbf{C} \in \text{co-CSP}(\mathbf{B}_{ps})$ because a terminal node is still accessible from the source nodes, and we also have that $\mathbf{C} \rightarrow \mathbf{A}$ because the reverse renaming of elements is a homomorphism from \mathbf{C} to \mathbf{A}' , and hence to \mathbf{A} .




Fig. 2. (1,3)-treewidth duality for the structure \mathbf{B}_{ps} .

2.3 Datalog and infinitary logics

For logical descriptions of the three dualities, we use first-order logic, the logic programming language Datalog, and its restriction, linear Datalog, and also some infinitary finite-variable logics. We assume that the reader is familiar with first-order logic, and we now briefly describe the basics of Datalog (for more details, see, e.g., [43]).

Fix a vocabulary τ . A *Datalog program* is a finite set of rules of the form $t_0 : - t_1, \dots, t_n$ where each t_i is an atomic formula $R(x_{i_1}, \dots, x_{i_k})$. Then t_0 is called the *head* of the rule, and the sequence t_1, \dots, t_n the *body* of the rule. The predicates occurring in the heads of the rules are not from τ and are called *IDBs* (from “intensional database predicates”), while all other predicates come from τ and are called *EDBs* (from “extensional database predicates”). One of the IDBs, which is usually 0-ary in our case, is designated as the *goal predicate* of the program. Since the IDBs may occur in the bodies of rules, each Datalog program is a recursive specification of the IDBs, with semantics obtained via least fixed-points of monotone operators. The goal predicate is assumed to be initially set to **false**, and we say that a Datalog program *accepts* a τ -structure \mathbf{A} if its goal predicate evaluates to **true** on \mathbf{A} .

For $0 \leq j \leq k$, a (j, k) -*Datalog program* is a Datalog program with at most j variables in the head and at most k variables per rule. A Datalog program is called *linear* if every rule in it has at most one occurrence of an IDB in its body.

A class \mathcal{C} of structures is said to be definable in (linear) (j, k) -Datalog if there is a (linear) (j, k) -Datalog program which accepts precisely the structures from \mathcal{C} .

Note that, for any Datalog program, the class \mathcal{C} of all structures accepted by the program is closed under extension (that is, if a structure \mathbf{A} has a substructure \mathbf{A}' which is in \mathcal{C} then \mathbf{A} is also in \mathcal{C}). Every class of the form $\text{co-CSP}(\mathbf{B})$ has this monotonicity property, but it is not the case for $\text{CSP}(\mathbf{B})$. Hence, when using Datalog to study CSPs, one usually speaks of definability of $\text{co-CSP}(\mathbf{B})$ in (some version of) Datalog.

Example 11. Consider the structure \mathbf{B}_{3H} from Example 5. It is well known that HORN 3-SAT can be solved by the unit propagation algorithm which can be represented as the following Datalog program.

$$\begin{aligned} T(X) &: - U(X) \\ T(Z) &: - P(X, Y, Z), T(X), T(Y) \\ \text{unsat} &: - N(X, Y, Z), T(X), T(Y), T(Z) \end{aligned}$$

Hence, $\text{co-CSP}(\mathbf{B}_{3H})$ is definable in (1,3)-Datalog.

Example 12. The following linear (2,4)-Datalog program accepts a graph (as a structure with one binary relation E) if and only if the graph is non-2-colourable

$$\begin{aligned} \text{Odd}(X, Y) &: - E(X, Y) \\ \text{Odd}(X, Y) &: - \text{Odd}(X, Z), E(Z, T), E(T, Y) \\ \text{non2col} &: - \text{Odd}(X, X) \end{aligned}$$

It is easy to see how to modify this program so that it accepts a digraph if and only if the digraph is not 2-colourable (just add all rules obtained from the second rule by permuting Z, T , and Y in the part $E(Z, T), E(T, Y)$). Hence, $\text{co-CSP}(\mathbf{K}_2)$ is definable in linear (2,4)-Datalog.

The notion of a *canonical* (j, k) -Datalog program for a τ -structure \mathbf{B} has proved to be useful in the study of dualities [26]. Let $\tau = \{R_1, \dots, R_n\}$, and let S_0, S_1, \dots, S_p be an enumeration of relations of arity j on B that can be expressed by a first-order $\exists\wedge$ -formula over \mathbf{B} . Assume that S_0 is the empty relation. For each S_i , introduce a j -ary IDB I_i . Then the canonical (j, k) -Datalog program for \mathbf{B} involves the IDBs I_0, \dots, I_p and EDBs R_1, \dots, R_n , and contains all the rules with at most k variables with the following property: if every I_i in the rule is replaced by S_i and every R_s by $R_s^{\mathbf{B}}$, then every assignment of elements of B to the variables that satisfies the conjunction of atomic formulas in the body must also satisfy the atomic formula in the head. Finally, introduce one 0-ary IDB G together with the rule $G : - I_0(x_1, \dots, x_j)$, and make G the goal predicate of the program. The *canonical linear* (j, k) -Datalog program for \mathbf{B} consists of all linear rules from the canonical program described above.

Our definitions of *infinitary logics* are inspired by [19, 43, 45]. Let $L_{\infty\omega}$ be the first-order logic extended with infinitary conjunctions \bigwedge and infinitary disjunctions \bigvee . For every $k \geq 0$, let $\exists L_{\infty\omega}^k$ be the existential positive (i.e., without

negation and universal quantifiers) fragment of $L_{\infty\omega}$ with at most k different variables. A (possibly infinitary) conjunction $\bigwedge \Phi$ of $L_{\infty\omega}$ -formulas is said to be *j-restricted* if every formula from Φ that contains more than j free variables is quantifier-free, and it is said to be *strongly j-restricted* if, in addition, at most one formula in Φ having quantifiers is not a sentence. Then $\exists L_{\infty\omega}^{j,k}$ is the fragment of $\exists L_{\infty\omega}^k$ obtained by using atomic formulas, existential quantification, arbitrary disjunctions, and *j-restricted* conjunctions. It is known that every class of structures definable in (j, k) -Datalog is also definable in $\exists L_{\infty\omega}^{j,k}$. The logic $\exists M_{\infty\omega}^{j,k}$ is defined similarly to $\exists L_{\infty\omega}^{j,k}$, but with strongly *j-restricted* conjunctions, and definability in linear (j, k) -Datalog implies definability in this logic.

We will also need the infinitary *counting* logics. Let $C_{\infty\omega}$ (see [3, 53]) be the logic whose formulas are obtained from atomic formulas by using negation, infinitary conjunction and disjunction, and counting quantifiers ($\exists^i x$ for any $i \geq 0$). The fragment $C_{\infty\omega}^k$ consists of those formulas of $C_{\infty\omega}$ in which at most k distinct variables appear, and $C_{\infty\omega}^\omega = \bigcup_{k \in \omega} C_{\infty\omega}^k$.

2.4 Pebble games

We will now define two pebble games, the pebble-relation game and the existential pebble game, which have proved to be very useful in the analysis of pathwidth and treewidth dualities. These games have been introduced in [19] and [44], respectively.

Let $0 \leq j \leq k$, and let \mathbf{A} and \mathbf{B} be τ -structures. The (j, k) -*pebble-relation* (or (j, k) -*PR*) *game* on (\mathbf{A}, \mathbf{B}) is played between two players, the *Spoiler* and the *Duplicator*. A configuration of the game consists of a subset $I \subseteq A$ with $|I| \leq k$ and a collection of partial homomorphisms $T \subseteq \text{hom}(\mathbf{A}_{|I}, \mathbf{B})$. If $T \subseteq \text{hom}(\mathbf{A}_{|I}, \mathbf{B})$ then we say that I is the *domain* of T . For a subset $J \subseteq I$, let $T|_J$ denote the set $\{f|_J \mid f \in T\}$.

Initially, $I = \emptyset$ and T contains the (unique) homomorphism from $\mathbf{A}_{|\emptyset}$ to \mathbf{B} . Each round of the game consists of a move of the Spoiler and a move of 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* 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} to be a non-empty subset of the domain I^n of T^n . The Duplicator responds by restricting every function in T^n onto 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 family $T^{n+1} \subseteq \text{hom}(\mathbf{A}_{|I^{n+1}}, \mathbf{B})$ such that $T|_{I^n}^{n+1} \subseteq T^n$.

The Spoiler wins the game if the response of the Duplicator sets T^{n+1} to \emptyset , i.e., the Duplicator cannot extend successfully any of the partial homomorphisms

from T^n . Otherwise, the game resumes. The Duplicator wins the game if he has a strategy that allows him to play “forever”, i.e., if the Spoiler can never win a round of the game. The notion of winning strategy for the Duplicator can be conveniently formalised as follows.

Definition 6. *Let $0 \leq j < k$, 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}, \mathbf{B}) if there is a non-empty family \mathcal{H} of sets of partial homomorphisms such that:*

1. *for every $T \in \mathcal{H}$, $T \subseteq \text{hom}(\mathbf{A}_{|I}, \mathbf{B})$ for some $I \subseteq A$, $|I| \leq k$, and $\emptyset \notin T$,*
2. *\mathcal{H} is closed under restrictions: for every $T \in \mathcal{H}$ with domain I and every $I' \subseteq I$, we have that $T_{|I'} \in \mathcal{H}$,*
3. *\mathcal{H} has the (j, k) -forth property: for every $T \in \mathcal{H}$ with domain I , $|I| \leq j$, and every superset I' of I with $|I'| \leq k$, there exists $T' \in \mathcal{H}$ with domain I' such that $T'_I \subseteq T$.*

The intuition behind the above definition is that every set T in a winning strategy corresponds to a winning configuration for the Duplicator in the game.

If we impose the restriction that every configuration in the (j, k) -PR game consists of a single function (i.e., in every round, the Duplicator commits to a particular partial homomorphism) then the obtained game is known as the *existential (j, k) -pebble game*. The notion of a winning strategy for the Duplicator in this game is obtained in a natural way from the one in Definition 6, by restricting each set T to consist of a single partial homomorphism.

Note that if we have a homomorphism $h : \mathbf{A} \rightarrow \mathbf{B}$ then the Duplicator always has a winning strategy in any PR or existential pebble game on (\mathbf{A}, \mathbf{B}) : to win, the Duplicator only has to always include the suitable restriction of the homomorphism h in his response. However, the converse does not always hold. That is, the existence of a winning strategy for the Duplicator on (\mathbf{A}, \mathbf{B}) does not, in general, imply that $\mathbf{A} \rightarrow \mathbf{B}$ (see Example 13 below). Thus, the structures \mathbf{B} , for which the converse also holds (for a particular type of game), must have some special properties. These properties are closely related with dualities, as we will discuss in Sections 4 and 5.

Example 13. Let \mathbf{A} be the undirected cycle with 5 nodes and \mathbf{B} the undirected cycle with 6 nodes. Obviously, we have $\mathbf{A} \not\rightarrow \mathbf{B}$, but the Duplicator still wins the existential $(1, 2)$ -pebble game. Indeed, fix any two adjacent elements, b_1 and b_2 in \mathbf{B} , and let the winning strategy simply contain all partial homomorphisms that have at most two-element domains and range $\{b_1, b_2\}$. It is straightforward to check that this is indeed a winning strategy. However, it is not hard to verify that the Spoiler wins the existential $(2, 3)$ -pebble game on (\mathbf{A}, \mathbf{B}) .

2.5 Algebraic background

The algebraic approach to constraint satisfaction (see, e.g., [11–13, 16, 48]) has proved to be extremely successful. It provides a convenient dual language to

analyse CSPs, and, more importantly, allows one to use powerful machinery from universal algebra.

First, let us formally define polymorphisms of relations and structures.

Definition 7. Let f be an n -ary operation on B , and R a relation on B . Then f is said to be a polymorphism of R (or R is invariant under f) if, for any tuples $\bar{a}_1, \dots, \bar{a}_n \in R$, the tuple obtained by applying f componentwise also belongs to the relation R .

An operation is called a polymorphism of a relational structure if it is a polymorphism of every relation in the structure. Let $\text{Pol}(\mathbf{B})$ denote the set of all polymorphisms (of all arities) of a structure \mathbf{B} .

For τ -structures $\mathbf{B}_1, \dots, \mathbf{B}_n$, define the *direct product* structure $\mathbf{C} = \prod_{i=1}^n \mathbf{B}_i$ to be a τ -structure with base set $C = B_1 \times \dots \times B_n$, and, for any m -ary $R \in \tau$, let $(\mathbf{a}_1, \dots, \mathbf{a}_m) \in R^{\mathbf{C}}$ if and only if $(\mathbf{a}_1[i], \dots, \mathbf{a}_m[i]) \in R^{\mathbf{B}_i}$ for each $1 \leq i \leq n$. As usual, the direct product of n copies of a structure \mathbf{B} is called the n -th power of \mathbf{B} , and is denoted \mathbf{B}^n . It is easy to check that the n -ary polymorphisms of \mathbf{B} are precisely the homomorphisms from \mathbf{B}^n to \mathbf{B} .

Example 14. It is straightforward to verify that the Boolean relation $OR = \{0, 1\}^2 \setminus \{(0, 0)\}$ is invariant under the binary operation \max on $\{0, 1\}$, but is not invariant under the operation \min .

One nice feature of the polymorphisms is that they allow one to simultaneously deal with structures over different vocabularies. For example, it is known (see [12] or [40]) that if τ_1 -structure \mathbf{B}_1 and τ_2 -structure \mathbf{B}_2 have the same universe and $\text{Pol}(\mathbf{B}_1) \subseteq \text{Pol}(\mathbf{B}_2)$ then every relation in \mathbf{B}_2 can be defined by a primitive positive first-order formula (i.e., $\exists\wedge$ -formula with equality) in \mathbf{B}_1 , and hence the problem $\text{CSP}(\mathbf{B}_2)$ is polynomial-time (even logarithmic-space) reducible to $\text{CSP}(\mathbf{B}_1)$. In particular, if $\text{Pol}(\mathbf{B}_1) = \text{Pol}(\mathbf{B}_2)$ then $\text{CSP}(\mathbf{B}_1)$ and $\text{CSP}(\mathbf{B}_2)$ are equivalent. Hence, it is very convenient to group relational structures according to their polymorphisms. Note that sets of operations of the form $\text{Pol}(\mathbf{B})$ are *clones* of operations, they are well-studied objects in universal algebra (see, e.g., [60]).

We will now define some types of operations which will be useful in the subsequent sections.

Definition 8. An n -ary operation f on B is called idempotent if it satisfies the identity $f(x, \dots, x) = x$.

- A binary commutative idempotent operation f is called a 2-semilattice operation if it satisfies the identity $f(x, f(x, y)) = f(x, y)$.
- An n -ary ($n \geq 2$) operation f is called totally symmetric if $f(x_1, \dots, x_n) = f(y_1, \dots, y_n)$ whenever $\{x_1, \dots, x_n\} = \{y_1, \dots, y_n\}$. If, in addition, f is idempotent then we say that it is a TSI operation.
- An n -ary ($n \geq 3$) operation is called an NU (near-unanimity) operation if it satisfies the identities

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

- A ternary NU operation is called a majority operation.
- An n -ary ($n \geq 2$) idempotent operation is called a WNU (weak NU) operation if it satisfies the identities

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

- Example 15.* 1. For any binary idempotent operation f , the following conditions are equivalent: (a) f is a TSI operation, (b) f is a WNU operation, and (c) f is commutative.
2. A binary operation g is called conservative if $g(a, b) \in \{a, b\}$ for all a, b . Any binary commutative conservative operation is a 2-semilattice operation.
3. Let f be a binary idempotent commutative associative operation. Then f is called a semilattice operation. It is easy to see that f is also a 2-semilattice operation, and, for any $n \geq 2$, the operation $f(x_1, f(x_2, f(\dots, f(x_{n-1}, x_n) \dots))$ is a TSI operation.
4. It is easy to check that the (ternary) median operation on a totally ordered set is a majority operation.
5. Any TSI operation and any NU operation is a WNU operation. Also, the Boolean affine operation $f(x, y, z) = x + y + z \pmod{2}$ is a WNU operation.

Example 16. Schaefer’s celebrated dichotomy theorem for Boolean CSP can be restated (see, e.g., [12, 16, 48]) as follows. For a Boolean core structure \mathbf{B} , if \mathbf{B} has a semilattice polymorphism, or a majority polymorphism, or the affine polymorphism, then $\text{CSP}(\mathbf{B})$ is in **PTIME**. In all other cases, $\text{CSP}(\mathbf{B})$ is **NP**-complete. A refinement of this theorem, including a classification for definability in Datalog and its restrictions, can be found in [50].

The subsequent definitions in this subsection are sketchy, for more details see the surveys [13, 15] or the monograph [38].

Definition 9. A finite algebra is a pair $\mathbb{A} = (A, F)$ where A is a finite set and $F = (f_i)_{i \in I}$ is a family of finitary operations on A . For a relational structure \mathbf{B} , the algebra $\mathbb{A}_{\mathbf{B}} = (B, \text{Pol}(\mathbf{B}))$ is called the algebra associated with \mathbf{B} .

Definition 10. A variety is a class of algebras closed under taking homomorphic images, subalgebras, and (possibly infinite) direct products. The variety generated by a finite algebra \mathbb{A} , denoted $\text{var}(\mathbb{A})$, consists of all homomorphic images of subalgebras of direct powers of \mathbb{A} .

Every finite algebra \mathbb{A} can be assigned a set of *types*. The types are numbers from **1** to **5**, and they correspond to different possible basic “local behaviours” of the algebra. The correspondence is as follows:

- type **1** – unary algebra,
- type **2** – vector space over a finite field,
- type **3** – 2-element Boolean algebra,
- type **4** – 2-element lattice,
- type **5** – 2-element semilattice.

A variety is said to *admit* a type \mathbf{i} if this type occurs in some finite algebra in the variety, and it *omits* type \mathbf{i} otherwise.

It is known (see [12, 13, 48]) that if, for a core structure \mathbf{B} , the variety $\text{var}(\mathbb{A}_{\mathbf{B}})$ admits type $\mathbf{1}$ then $\text{CSP}(\mathbf{B})$ is **NP**-complete. Moreover, all core structures \mathbf{B} that are known to give rise to **NP**-complete problems $\text{CSP}(\mathbf{B})$ do satisfy this condition. It has been conjectured that all other core structures give rise to problems in **PTIME**, and this conjecture has been confirmed in many important cases (see, e.g., [7, 10, 12, 13]). For other results about the correspondence between the type set of $\text{var}(\mathbb{A}_{\mathbf{B}})$ on one side and the computational and descriptive complexity of $\text{CSP}(\mathbf{B})$ on the other side, see [3, 13, 50].

3 Finite duality

Arguably, the simplest case of duality is that of finite duality. In this section, we outline several characterisations of constraint satisfaction problems with this property. We shall address, in particular, questions about the relationship of finite duality to definability in first-order logic (FO), the nature of the obstruction set of a structure with finite duality, and the (meta-)problem of recognising such structures.

Recall from Example 8 that the transitive tournament \mathbf{T}_n has an obstruction set consisting of a single structure \mathbf{P}_n . In general, a structure with finite duality might not have a set of obstructions that consists of a single structure:

Example 17. Let $\mathbf{B} = \langle \{0, 1\}; R, \{0\}, \{1\} \rangle$ where $R = \{(0, 0), (0, 1), (1, 0)\}$. Viewing structures of this type as coloured digraphs (with colours given by the unary relations), it is easy to see that $\mathbf{A} \not\rightarrow \mathbf{B}$ if and only if there exists a vertex v of \mathbf{A} which is coloured with both colours 0 and 1, or an edge (a, b) with both endpoints coloured 1. Consequently \mathbf{B} has a two-element obstruction set, one structure \mathbf{A}_1 consisting of a single vertex with two colours, the other structure \mathbf{A}_2 consisting of one directed edge with both ends coloured 1. It is easy to see that \mathbf{B} does not have a one-element obstruction set.

Example 18. Recall the problem UNREACHABILITY, or $\text{CSP}(\mathbf{B}_{unr})$ from Example 3. It is not difficult to see that \mathbf{B}_{unr} does not have finite duality. As in Example 17, we can view structures as coloured digraphs. Note that any path with ends coloured 0 and 1 does not have a homomorphism to \mathbf{B}_{unr} , but any proper substructure of the path does. If \mathcal{O} is a finite obstruction set for \mathbf{B}_{unr} , then one can find a long enough path \mathbf{P} (with coloured ends) such that every structure in \mathcal{O} can have only non-surjective homomorphisms (if any) to \mathbf{P} . Hence, either none of the structures in \mathcal{O} has a homomorphism to \mathbf{P} or some structure in \mathcal{O} has a homomorphism to \mathbf{B}_{unr} . In either case, \mathcal{O} cannot be an obstruction set for \mathbf{B}_{unr} .

It is easy to see that if a structure \mathbf{B} has finite duality then $\text{CSP}(\mathbf{B})$ is FO-definable; in fact, $\text{co-CSP}(\mathbf{B})$ is definable in existential positive FO, or said differently, it is definable in Datalog without IDBs other than the goal predicate.

Indeed, let \mathbf{C} be a τ -structure with $C = \{c_1, \dots, c_l\}$, and consider the following sentence $T_{\mathbf{C}} = \exists x_1 \dots \exists x_l \bigwedge_{R \in \tau} \bigwedge_{(c_{i_1}, \dots, c_{i_r}) \in R^{\mathbf{C}}} R(x_{i_1}, \dots, x_{i_r})$. It is well known and easy to check that, for any τ -structure \mathbf{A} , we have $\mathbf{C} \rightarrow \mathbf{A}$ if and only if \mathbf{A} satisfies $T_{\mathbf{C}}$. Hence, if \mathcal{O} is a finite obstruction set for a τ -structure \mathbf{B} , then a τ -structure \mathbf{A} belongs to $\text{co-CSP}(\mathbf{B})$ if and only if the sentence $\bigvee_{\mathbf{C} \in \mathcal{O}} T_{\mathbf{C}}$ holds true in \mathbf{A} .

Atserias ([2], see also [59]) has shown that the converse also holds: if $\text{CSP}(\mathbf{B})$ is FO-definable then \mathbf{B} has finite duality. We now show how this result follows from other, more recent, results.

The next result is stated in [50] in a slightly weaker form, but in fact the proof relies on a characterisation of structures with finite duality from [49], and immediately yields the following:

Theorem 1 ([50]). *If a structure \mathbf{B} does not have finite duality then $\text{CSP}(\mathbf{B})$ is **LOGSPACE**-hard under first-order reductions.*

Recall that the complexity class **non-uniform AC**⁰ consists of all languages accepted by polynomial-size constant-depth families of Boolean circuits (see, e.g., [53]). It is known that any FO-definable class of structures belongs to this complexity class (see Theorem 6.4 of [53]). Moreover, any problem which is **LOGSPACE**-hard under first-order reductions cannot lie in **non-uniform AC**⁰ because there are problems in **LOGSPACE** which are not in **non-uniform AC**⁰ (see [28]) and **non-uniform AC**⁰ is closed under first-order reductions. These facts and Theorem 1 imply the following result.

Theorem 2. *For any structure \mathbf{B} , the following conditions are equivalent:*

1. \mathbf{B} has finite duality.
2. $\text{CSP}(\mathbf{B})$ is FO-definable.
3. $\text{CSP}(\mathbf{B})$ is in **non-uniform AC**⁰.

Let us now consider the question about the nature of finite obstruction sets.

Definition 11. *Let \mathbf{A} be a τ -structure. The incidence multigraph of \mathbf{A} , denoted $\text{Inc}(\mathbf{A})$, is defined as the bipartite multigraph with parts A and $\text{Block}(\mathbf{A})$, where $\text{Block}(\mathbf{A})$ consists of all pairs (R, \bar{a}) such that $R \in \tau$ and $\bar{a} \in R^{\mathbf{A}}$, and with edges $e_{a,i,Z}$ joining $a \in A$ to $Z = (R, (a_1, \dots, a_r)) \in \text{Block}(\mathbf{A})$ when $a_i = a$. We say that the structure \mathbf{A} is a τ -tree (or simply a tree) if its incidence multigraph is a tree (in particular, it has no multiple edges).*

Theorem 3 ([56, 57]). *If a finite structure has finite duality, then it admits an obstruction set consisting of finitely many trees. Conversely, for any finite set \mathcal{O} of trees, there is a structure \mathbf{B} that can be explicitly constructed from \mathcal{O} such that $\mathcal{O} = \mathcal{O}_{\mathbf{B}}$.*

Note that the structure \mathbf{B} obtained in the above theorem may not be a core; in fact, it may be much larger than its core.

We now give an algebraic characterisation of structures with finite duality.

Definition 12. Let R be a relation on the set A . An n -ary operation f on A is a 1-tolerant polymorphism of R if, for any tuples $\bar{a}_1, \dots, \bar{a}_n$ at least $n - 1$ of which belong to R , the tuple obtained by applying f componentwise also belongs to R .

Theorem 4 ([49]). A structure \mathbf{B} has finite duality if and only if its core has a 1-tolerant NU polymorphism.

In fact, the arity of such a 1-tolerant NU polymorphism is determined by the total number of tuples in the relations of minimal obstructions. A structure \mathbf{A} is a *critical obstruction* of \mathbf{B} if $\mathbf{A} \not\rightarrow \mathbf{B}$ and $\mathbf{A}' \rightarrow \mathbf{B}$ for any proper substructure \mathbf{A}' of \mathbf{A} . Call any tuple of any relation of a structure \mathbf{A} a *hyperedge* of this structure. Then we have the following:

Theorem 5 ([49]). The core of \mathbf{B} admits a 1-tolerant NU polymorphism of arity $n + 1$ if and only if each critical obstruction of \mathbf{B} has at most n hyperedges.

Example 19. 1. The transitive tournament \mathbf{T}_n of Example 8 admits a 1-tolerant NU polymorphism of arity $n + 1$, but none of smaller arity (even though it has a majority polymorphism).
 2. The structure \mathbf{B} of Example 17 admits a 1-tolerant NU polymorphism of arity 4, but not 3. Indeed, if m was a ternary 1-tolerant NU polymorphism of the relation R of \mathbf{B} , we would have $(1, 1) = (m(1, 1, 0), m(1, 0, 1)) \in R$, which is false. On the other hand, it is straightforward to check that the 4-ary operation f such that $f(x_1, \dots, x_4) = 1$ if and only if at most one x_i is equal to 0 is a 1-tolerant NU polymorphism of \mathbf{B} .

Call a relation R on B *biredundant* if the projection of R onto some two coordinates is the equality relation on some subset $C \subseteq B$ with $|C| \geq 2$.

Theorem 6 ([23, 49]). Let \mathbf{B}_1 and \mathbf{B}_2 be structures such that \mathbf{B}_1 is a core with finite duality and $\text{Pol}(\mathbf{B}_1) \subseteq \text{Pol}(\mathbf{B}_2)$. Then the following holds.

1. If \mathbf{B}_2 does not have finite duality then $\text{CSP}(\mathbf{B}_2)$ is **LOGSPACE**-complete.
2. If none of the relations in \mathbf{B}_2 is biredundant then \mathbf{B}_2 also has finite duality. If \mathbf{B}_2 is a core then the converse holds as well.

Example 20. We will now describe Boolean structures that are cores with finite duality. (Boolean non-core structures trivially have this property). It can be derived from [49, 50] that these are precisely the (Boolean) core structures \mathbf{B} without biredundant relations and such that (at least) one of the ternary operations $x \vee (y \wedge \bar{z})$ and $x \wedge (y \vee \bar{z})$ is a polymorphism of \mathbf{B} .

We shall now describe a simple algorithm to determine if a structure \mathbf{B} has finite duality. A slight modification of this algorithm also provides a way of producing solutions of a CSP with finite duality. First, we require a few straightforward definitions.

Definition 13. Let \mathbf{A} be a structure and let $a, b \in A$. We say that the element a dominates the element b if, in any tuple \bar{t} in any relation R in \mathbf{A} , replacement in \bar{t} of any number of occurrences of b by a yields a tuple also in R .

For example, if a dominates b and $(b, c, b) \in R$ then (a, c, b) , (b, c, a) , and (a, c, a) are all in R . Note that this notion is a direct generalisation of the notion of domination in graph theory.

Recall from Section 2.5 the definition of the n -th power of a structure. Obviously, the second power \mathbf{B}^2 of a structure \mathbf{B} is called the *square* of \mathbf{B} . The *diagonal* $\Delta(\mathbf{B}^2)$ of the square \mathbf{B}^2 is the substructure of \mathbf{B}^2 induced by the set $\{(b, b) \mid b \in B\}$. Note that $\Delta(\mathbf{B}^2)$ is isomorphic to \mathbf{B} .

Definition 14. A structure \mathbf{A} is said to dismantle to its substructure \mathbf{C} if there exists a sequence of induced substructures $\mathbf{A}_0, \dots, \mathbf{A}_k$ of \mathbf{A} such that (i) $\mathbf{A}_0 = \mathbf{A}$, (ii) $\mathbf{A}_k = \mathbf{C}$ and (iii) for each $0 \leq j < k$ the structure \mathbf{A}_{j+1} is obtained from \mathbf{A}_j by removal of a dominated element of \mathbf{A}_j .

It is known [49] that the procedure of dismantling can always be done greedily, by successively removing arbitrary dominated elements in substructures of \mathbf{A} to eventually obtain \mathbf{C} .

Theorem 7 ([49]). A structure \mathbf{B} has finite duality if and only if it has a retract \mathbf{A} whose square \mathbf{A}^2 dismantles to its diagonal $\Delta(\mathbf{A}^2)$.

Example 21. Consider the tournament \mathbf{T}_3 (see Example 8). We know that \mathbf{T}_3 is a core with finite duality, so its square \mathbf{T}_3^2 should dismantle to its diagonal. We will now show that this is indeed the case. The process of dismantling is shown on Fig. 3.

The digraph \mathbf{T}_3^2 is shown in Fig. 3, top-left. The vertices $(2, 0)$ and $(0, 2)$ are dominated by all vertices, so they are removed, and the resulting digraph is shown in Fig. 3, top-right. Next, the vertices $(1, 0)$ and $(0, 1)$ are now dominated by $(0, 0)$, so they are removed (see Fig. 3, bottom-left). Finally, the vertices $(1, 2)$ and $(2, 1)$ are now dominated by $(2, 2)$, so they are removed as well, which leaves only the diagonal $\Delta(\mathbf{T}_3^2)$, shown in Fig. 3, bottom-right.

From Theorem 7, the problem of recognising structures with finite duality is in **NP**. Indeed, one only needs to guess a mapping ϕ from B onto its subset A , and then to check that the induced (by A) substructure \mathbf{A} of \mathbf{B} is a retract of \mathbf{B} (via ϕ), then to form the square \mathbf{A}^2 and, finally, to check (greedily) that \mathbf{A}^2 dismantles to its diagonal, which clearly can all be done in polynomial time.

Theorem 8 ([49]).

1. The problem of deciding whether a given structure \mathbf{B} has finite duality is **NP**-complete.
2. The problem of deciding whether a given structure \mathbf{B} is a core with finite duality is in **PTIME**.

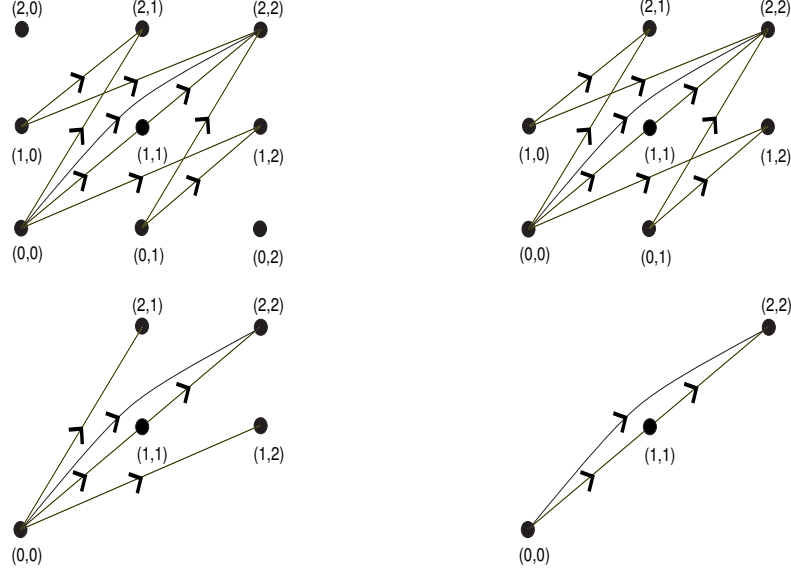


Fig. 3. Dismantling T_3^2 to its diagonal.

We now present a slight modification of this algorithm which will yield a solution to the CSP when one exists. In a product $\mathbf{A} \times \mathbf{B}$, an element (a, b) is said to be *dominated in the second coordinate* if it is dominated by an element of the form (a, b') . We say that $\mathbf{A} \times \mathbf{B}$ *dismantles in the second coordinate* to its substructure \mathbf{C} if \mathbf{C} can be obtained from $\mathbf{A} \times \mathbf{B}$ by successively removing elements that are dominated in the second coordinate.

Theorem 9 ([49]). *Let \mathbf{B} be a core with finite duality and let \mathbf{A} be a structure similar to \mathbf{B} . Let \mathbf{C} be a structure with no dominations which is obtained from $\mathbf{A} \times \mathbf{B}$ by dismantling in the second coordinate. Then $\mathbf{A} \rightarrow \mathbf{B}$ if and only if \mathbf{C} is the graph of a homomorphism from \mathbf{A} to \mathbf{B} .*

In other words, the procedure is as follows: (i) dismantle greedily the product $\mathbf{A} \times \mathbf{B}$ in the second coordinate until no dominations are left; (ii) check if the resulting set is of the form $C = \{(a, \phi(a)) : a \in A\}$ for some map $\phi : A \rightarrow B$; if it is, verify that ϕ is a homomorphism. Then either it is and ϕ is the desired solution, or else there is no homomorphism from \mathbf{A} to \mathbf{B} . Note that the result remains valid not only for cores, but for any structure \mathbf{B} whose square dismantles to the diagonal.

4 Bounded pathwidth duality

In this section we consider bounded pathwidth duality, which is a property shared by all structures \mathbf{B} such that $\text{CSP}(\mathbf{B})$ is currently known to belong to **NL**. The following result ties together pathwidth dualities, linear Datalog, and PR games.

Theorem 10 ([19]). *For any structure \mathbf{B} , the following conditions are equivalent:*

1. \mathbf{B} has (j, k) -pathwidth duality.
2. $\text{co-CSP}(\mathbf{B})$ is definable in linear (j, k) -Datalog.
3. $\text{co-CSP}(\mathbf{B})$ is definable by the canonical linear (j, k) -Datalog program for \mathbf{B} .
4. $\text{co-CSP}(\mathbf{B})$ is definable in $\exists M_{\infty\omega}^{j,k}$.
5. $\text{CSP}(\mathbf{B})$ is the class of all structures \mathbf{A} such that the Duplicator wins the (j, k) -PR game on (\mathbf{A}, \mathbf{B}) .

If these conditions hold then $\text{CSP}(\mathbf{B})$ is in NL.

Dalmau [19] also provides other equivalent conditions, including definability in other infinitary finite-variable logics and in fragments of second-order logic.

We will now give examples of structures with bounded pathwidth duality.

Example 22. An *oriented path* is a digraph obtained from a path by orienting its edges in some way. A digraph is called a *local tournament* if the set of out-neighbours of any vertex induces a tournament. For example, all transitive tournaments and all directed paths (see Example 8) are local tournaments. It was shown in [35, 36] that any digraph \mathbf{H} that is an oriented path or an acyclic local tournament has an obstruction set consisting of oriented paths. Since any oriented path has pathwidth at most $(1, 2)$, it follows that \mathbf{H} has $(1, 2)$ -pathwidth duality.

Example 23. An *oriented cycle* is a digraph obtained from a cycle by orienting its edges in some way. An oriented cycle is called *balanced* if it has the same number of edges in one direction and in the other, and it is *unbalanced* otherwise. It was shown in [37] that any unbalanced oriented cycle \mathbf{H} has an obstruction set consisting of oriented paths and oriented cycles. Since oriented cycles have pathwidth at most $(2, 3)$ (see Example 7), such a digraph \mathbf{H} has $(2, 3)$ -pathwidth duality. Moreover, if the difference between the number of edges in \mathbf{H} going in one direction and the number of edges in the other direction is exactly one then \mathbf{H} has an obstruction set consisting only of oriented paths [37], and so it has $(1, 2)$ -pathwidth duality.

Example 24. A binary relation on B is called *implicational* (or *0/1/all*) if it has one of the following three forms: (1) $C \times D$ for some $C, D \subseteq B$, (2) $\{(c, f(c)) \mid c \in C\}$ for some $C \subseteq B$ and some permutation f on B , (3) $(\{c\} \times D) \cup (C \times \{d\})$ for some $C, D \subseteq B$, $c \in C$, and $d \in D$. A structure is called *implicational* if all of its relations are such. For example, it is easy to show (or see [19]) that the 2-SAT problem can be represented as $\text{CSP}(\mathbf{B})$ for an implicational structure \mathbf{B} (with universe $\{0, 1\}$). It was shown in [19] that every implicational structure has $(2, 3)$ -pathwidth duality.

Example 25. The class of *implicative hitting-set bounded* (IHS-B) relations was introduced in [18]. For $k \geq 2$, a Boolean relation is in k -IHS-B+ if it can be expressed as a CNF where each clause is of the form $\neg x$, $\neg x \vee y$, or $x_1 \vee \dots \vee x_k$.

Dually, a Boolean relation is in k -IHS-B− if it can be expressed as a CNF where each clause is of the form x , $\neg x \vee y$, or $\neg x_1 \vee \dots \vee \neg x_k$. It was shown in [19] that any structure \mathbf{B}_{ihs} (with universe $\{0, 1\}$) all whose relations are in k -IHS-B+ (or in k -IHS-B−) has $(k, k - 1 + \text{maxar}(\mathbf{B}_{ihs}))$ -pathwidth duality.

We mentioned in Section 2.5 that the polymorphisms of a structure \mathbf{B} determine the complexity of $\text{CSP}(\mathbf{B})$. Similarly, the polymorphisms determine whether a structure has bounded pathwidth duality.

Theorem 11 ([23, 50]). *Let \mathbf{B}_1 and \mathbf{B}_2 be relational structures with the same universe and such that $\text{Pol}(\mathbf{B}_1) \subseteq \text{Pol}(\mathbf{B}_2)$. If $\text{co-CSP}(\mathbf{B}_1)$ is definable in linear Datalog, then so is $\text{co-CSP}(\mathbf{B}_2)$.*

For a structure \mathbf{B} , let \mathbf{B}_c denote the structure obtained from \mathbf{B} by adding all elements of B as singleton unary relations.

Theorem 12 ([50]). *For a core structure \mathbf{B} , $\text{co-CSP}(\mathbf{B})$ is definable in linear Datalog if and only if $\text{co-CSP}(\mathbf{B}_c)$ is.*

Note that the polymorphisms of the structure \mathbf{B}_c in the above theorem are the idempotent polymorphisms of \mathbf{B} . Hence, for core structures, the idempotent polymorphisms determine whether a structure has bounded pathwidth duality.

The only currently known sufficient algebraic condition for general structures to have bounded pathwidth duality is given by the following result:

Theorem 13 ([21]). *If $|B| = k$ and \mathbf{B} has a majority polymorphism then \mathbf{B} has $(3k + 2, 3k + \text{maxar}(\mathbf{B}))$ -pathwidth duality.*

Note that Theorem 13 can be used to obtain bounded pathwidth duality for all structures from Examples 22-24 (though, with worse bounds). For example, it was shown in [24] that any oriented path and any unbalanced oriented cycle has a majority polymorphism, and the same can be shown for any acyclic local tournament. If \mathbf{B} is an implicational structure then, as shown in [17], \mathbf{B} has a majority polymorphism of a very specific form, the so-called dual discriminator.

For certain types of structures \mathbf{B} , the presence of a majority polymorphism is the dividing line, for $\text{CSP}(\mathbf{B})$, between membership in **PTIME** (which, by Theorems 10 and 13, becomes membership in **NL**) and **NP**-completeness, that is, either \mathbf{B} has a majority polymorphism or else $\text{CSP}(\mathbf{B})$ is **NP**-complete. For example, this is the case when \mathbf{B} is a structure \mathbf{B}_{lhc} (from Example 2) whose underlying digraph \mathbf{H} is undirected [25]. A combinatorial description of the boundary for the classification can also be found in [25].

The structures \mathbf{B}_{ihs} from Example 25, with $k \geq 3$ have bounded pathwidth duality, but do not have a majority polymorphism. However these structures are known to have an NU polymorphism of arity $k + 1$. Furthermore, it follows from known algebraic results (see, e.g., [60]) that a Boolean core structure \mathbf{B} has an NU polymorphism (of some arity) if and only if \mathbf{B} has a majority polymorphism or it is a structure of the form \mathbf{B}_{ihs} . Moreover, it can be derived from [50] that in all other cases \mathbf{B} does not have bounded pathwidth duality. That is, we obtain the following result:

Theorem 14. *Let $|B| = 2$. Then \mathbf{B} has bounded pathwidth duality if and only if \mathbf{B} has an NU polymorphism of some arity.*

It is not known whether the presence of an NU polymorphism is a sufficient condition for general structures to have bounded pathwidth duality. However, it is known that, in general, this condition is not necessary, i.e., there exist structures \mathbf{B} such that \mathbf{B} has bounded pathwidth duality, but no NU polymorphism of any arity. The simplest (known) structure \mathbf{B} with these properties is obtained as follows. Take the poset \mathbf{Q} whose Hasse diagram is shown in Fig. 4. Then \mathbf{B} is obtained from this poset by adding all elements of the universe as singleton unary relations. An explicit description of the minimal obstructions for \mathbf{B} in which the binary relation is a partial order (i.e., the so-called \mathbf{Q} -zigzags) can be found in [62]. It can be easily derived from this description that \mathbf{B} has the required properties.




Fig. 4. A poset without NU polymorphisms.

Let us now consider the question of which structures do not have bounded pathwidth duality. Trivially, any structure without bounded treewidth duality cannot have bounded pathwidth duality.

Example 26. Reconsider the structure \mathbf{B}_{ps} from Example 4. This structure has 1-treewidth duality, as shown in Example 10. By using the game technique (see below), it can be shown that \mathbf{B}_{ps} does not have bounded pathwidth duality (an alternative proof of this can be found in [1]).

At present, the most general algebraic necessary condition for the presence of bounded pathwidth duality is given by the following result.

Theorem 15 ([50]). *If a core structure \mathbf{B} has bounded pathwidth duality then the variety $\text{var}(\mathbb{A}_{\mathbf{B}})$ omits types **1**, **2**, and **5**.*

By using Theorem 9.11 of [38], the previous theorem can be re-stated as follows.

Theorem 16. *If a core structure \mathbf{B} has bounded pathwidth duality then \mathbf{B} has ternary polymorphisms d_0, \dots, d_n , $n \geq 2$, satisfying the following identities:*

$$\begin{aligned} d_0(x, y, z) &= x, \\ d_n(x, y, z) &= z, \\ d_i(x, y, x) &= d_{i+1}(x, y, x) \text{ for all even } i < n, \\ d_i(x, y, y) &= d_{i+1}(x, y, y) \text{ for all even } i < n, \\ d_i(x, x, y) &= d_{i+1}(x, x, y) \text{ for all odd } i < n. \end{aligned}$$

Moreover, if a core structure \mathbf{B} does not have the polymorphisms described above (or, equivalently, the variety $\text{var}(\mathbf{A}_{\mathbf{B}})$ admits at least one of the types **1, 2**, and **5**) then $\text{CSP}(\mathbf{B})$ is hard for **PTIME** or for **Mod_pL** (for some prime p) under first-order reductions [50], and thus is unlikely to belong to **NL**.

One very natural question about pathwidth dualities is whether they form a proper hierarchy or the hierarchy collapses to some level. That is, the question is whether there exists a number j such that, for any $j' \geq j$, every structure with j' -pathwidth duality also has j -pathwidth duality.

It follows from Theorem 10 that, in order to prove that a certain structure \mathbf{B} does not have (j, k) -pathwidth duality, one only needs to provide a structure \mathbf{A} such that the Duplicator has a winning strategy in the (j, k) -PR game on (\mathbf{A}, \mathbf{B}) , but it holds that $\mathbf{A} \not\rightarrow \mathbf{B}$. This game technique was used in [21] to give a negative answer to the above question. Let us now describe the structures that were used in [21].

Let $n \geq 1$ and let \mathbf{B}_n be the structure with universe B_n and relations R_n^l , $1 \leq l \leq n$, defined as follows. The universe B_n is the set $\{1, \dots, n\} \times \{1, 2\}$. For every $1 \leq l \leq n$, R_n^l is a binary symmetric relation on B_n that consists of all pairs $((i, m), (i', m'))$ satisfying at least one of the following conditions:

- $i > l$, $i = i'$, $m = m'$
- $i = i' = l$, $m \neq m'$,
- $i < l$ and $i' \leq l$,
- $i \leq l$ and $i' < l$.

Theorem 17 ([21]). *For every $n \geq 2$, the structure \mathbf{B}_n does not have n -pathwidth duality, but \mathbf{B}_n has a majority polymorphism (and hence $(6n + 2)$ -pathwidth duality).*

The only known fact concerning the meta-problem for bounded pathwidth duality is that, for any fixed $k \geq 1$, the problem of recognising structures with $(1, k)$ -pathwidth duality is decidable [19].

5 Bounded treewidth duality

In this section we consider bounded treewidth duality. Arguably, it is the most important duality because it is one of the two most general basic properties of

relational structures \mathbf{B} that are known to guarantee that the problem $\text{CSP}(\mathbf{B})$ is solvable in polynomial time (the other property can be found, e.g., in [39]), and the vast majority of such structures have bounded treewidth duality. The notion of bounded treewidth duality has strong links with methods of solving constraint satisfaction problems based on establishing local consistency (see, e.g., [16, 22, 32, 41]).

The following result links together treewidth dualities, Datalog, infinitary logics, and existential pebble games.

Theorem 18 ([26, 45]). *For any structure \mathbf{B} , the following conditions are equivalent:*

1. \mathbf{B} has (j, k) -treewidth duality.
2. $\text{co-CSP}(\mathbf{B})$ is definable in (j, k) -Datalog.
3. $\text{co-CSP}(\mathbf{B})$ is definable by the canonical (j, k) -Datalog program for \mathbf{B} .
4. $\text{co-CSP}(\mathbf{B})$ is expressible in $\exists L_{\infty\omega}^{j,k}$.
5. $\text{CSP}(\mathbf{B})$ consists of all structures \mathbf{A} such that Duplicator has a winning strategy in the existential (j, k) -pebble game.

*If these conditions hold then $\text{CSP}(\mathbf{B})$ is in **PTIME**.*

To prevent possible confusion, we note that the paper [45] speaks about definability of $\text{co-CSP}(\mathbf{B})$ in k -Datalog meaning (k, k) -Datalog (in our notation). Hence, this does not exactly correspond to k -treewidth duality in our sense.

In [26] and some subsequent papers (e.g., in [52]), problems $\text{CSP}(\mathbf{B})$ that have j -treewidth duality (or bounded treewidth duality) are called *width- j* (or *bounded width*, respectively) problems.

We will now give some examples of structures with and without bounded treewidth duality. Note that, trivially, every structure with bounded pathwidth duality also has bounded treewidth duality.

Example 27. Recall the structure \mathbf{B}_{3H} from Example 5; the problem $\text{CSP}(\mathbf{B}_{3H})$ is precisely HORN-3-SAT. It follows from Example 11 that \mathbf{B}_{3H} has $(1, 3)$ -treewidth duality. By replacing relations $P^{\mathbf{B}_{3H}}$ and $N^{\mathbf{B}_{3H}}$ in \mathbf{B}_{3H} with k -ary relations $P^{\mathbf{B}_{kH}} = \{0, 1\}^k \setminus \{(1, \dots, 1, 1)\}$ and $N^{\mathbf{B}_{kH}} = \{0, 1\}^k \setminus \{(1, \dots, 1, 0)\}$, respectively, one obtains a structure \mathbf{B}_{kH} such that $\text{CSP}(\mathbf{B}_{kH})$ is exactly HORN- k -SAT. An obvious modification of Example 11 shows that \mathbf{B}_{kH} has $(1, k)$ -treewidth duality.

Example 28. Recall from Example 23 that each unbalanced oriented cycle has $(2, 3)$ -pathwidth duality. It was shown in [24] that if \mathbf{H} is a balanced oriented cycle then either it has bounded treewidth duality or $\text{CSP}(\mathbf{H})$ is **NP**-complete. The description of the boundary between the cases is rather involved.

Example 29. Let τ be the vocabulary (P, S, T, E) where P is ternary, E is binary, and S and T are unary relation symbols. Consider the τ -structure \mathbf{B} with 4-element universe $\{0, 1, a, b\}$ and relations defined as follows. Reconsider structures \mathbf{B}_{ps} from Example 4 and \mathbf{K}_2 from Example 9, and assume that the universe

of \mathbf{K}_2 is $\{a, b\}$. If $R \in \{P, S, T\}$ then define $R^{\mathbf{B}} = R^{\mathbf{B}_{ps}}$, and let $E^{\mathbf{B}} = E^{\mathbf{K}_2}$. We claim that the structure \mathbf{B} has 2-treewidth duality, but neither 1-treewidth duality nor bounded pathwidth duality. It is easy to see that if a connected τ -structure \mathbf{A} homomorphically maps to \mathbf{B} then either $E^{\mathbf{A}}$ is empty or else the other three relations in \mathbf{A} are empty. Hence, \mathbf{B} has an obstruction set consisting of structures from $\mathcal{O}_{\mathbf{B}_{ps}}$ and $\mathcal{O}_{\mathbf{K}_2}$ (suitably expanded with empty relations) and of finitely many structures in which some element appears both in the binary relation and in one of the other three relations. Since \mathbf{B}_{ps} has 1-treewidth duality and \mathbf{K}_2 has (2,3)-pathwidth duality, we conclude that \mathbf{B} has 2-treewidth duality. On the other hand, \mathbf{B}_{ps} does not have bounded pathwidth duality (see Example 26) and it is straightforward to show that \mathbf{K}_2 does not have 1-treewidth duality. Hence, \mathbf{B} cannot have either of these two properties.

Example 30. As we know from Example 1, the **H-COLORING** problem coincides with the problem $\text{CSP}(\mathbf{B}_{hc})$ where \mathbf{B}_{hc} is the (di)graph \mathbf{H} . If \mathbf{H} is a bipartite graph, then $\text{core}(\mathbf{H}) = \mathbf{K}_2$ and $\text{CSP}(\mathbf{B}_{hc})$ coincides with the 2-COLOURABILITY problem, and \mathbf{B}_{hc} hence has (2,3)-pathwidth duality (see Example 9). If \mathbf{H} is a non-bipartite graph then $\text{CSP}(\mathbf{B}_{hc})$ is **NP**-complete [31]. It is known (see, e.g., [32]) that in this case \mathbf{B}_{hc} does not have bounded treewidth duality (without any complexity-theoretic assumptions).

Example 31. A *triad* is a digraph obtained from three oriented paths by choosing one end of each path and identifying these three vertices. It is shown in [2] that there exists a triad \mathbf{H} such that $\text{CSP}(\mathbf{B}_{hc})$ is in **PTIME**, but \mathbf{H} does not have bounded treewidth duality.

Within the algebraic approach to the CSP, a different concept, *relational width*, is often very useful, see, e.g., [10]. This concept is applicable even for infinite sets of relations, but in the case of relational structures (with finite vocabulary) relational width is strongly related to treewidth duality, as we shall now see. Relational width is usually (e.g., in [7, 9, 10, 13]) defined using the “variable-value” form (as given in Section 1) of the constraint satisfaction problem. A straightforward translation into the homomorphism form goes as follows.

Definition 15. For $k \geq 1$, a family $\mathcal{M} = \{H_I \mid I \subseteq A, |I| \leq k\}$, where each H_I is a non-empty set of mappings from I to B , is called a *k-minimal family* for (\mathbf{A}, \mathbf{B}) if

1. for any $I' \subseteq I \subseteq A$, $|I| \leq k$, we have $H_{I'} = \{h|_{I'} \mid h \in H_I\}$, and
2. for any $I \subseteq A$ with $|I| \leq k$, any $h \in H_I$, any (n -ary) $R \in \tau$, and any $(a_1, \dots, a_n) \in R^{\mathbf{A}}$, there exists a tuple $(b_1, \dots, b_n) \in R^{\mathbf{B}}$ such that
 - (a) $h(a_i) = b_i$ for all $a_i \in I$,
 - (b) for any $J \subseteq A$ with $|J| \leq k$, there exists $h' \in H_J$ such that $h'(a_i) = b_i$ for all $a_i \in J$.

A structure \mathbf{B} is said to have relational width k if, for any structure \mathbf{A} such that there is a k -minimal family for (\mathbf{A}, \mathbf{B}) , we have $\mathbf{A} \rightarrow \mathbf{B}$. A structure \mathbf{B} has bounded relational width if it has relational width k for some k .

Note that property 2(a) shows that every member of every $H_I \in \mathcal{M}$ is a projective homomorphism, while properties 1 and 2(b) show that there is strong compatibility between different sets in \mathcal{M} . The following theorem links relational width and treewidth dualities.

Theorem 19. *For any structure \mathbf{B} , the following holds.*

1. *If \mathbf{B} has (j, k) -treewidth duality, then it has relational width k .*
2. *If \mathbf{B} has relational width k and $\text{maxar}(\mathbf{B}) = r$, then it has (k, k') -treewidth duality where $k' = \max(k, r)$.*

In particular, \mathbf{B} has bounded treewidth duality if and only if it has bounded relational width.

Proof. (1) Assume that \mathbf{B} has (j, k) -treewidth duality. Take an arbitrary structure \mathbf{A} such that there exists a k -minimal family $\mathcal{M} = \{H_I \mid I \subseteq A, |I| \leq k\}$ for (\mathbf{A}, \mathbf{B}) . We need to show that $\mathbf{A} \rightarrow \mathbf{B}$. By assumption and by Theorem 18, it is enough to exhibit a winning strategy \mathcal{H} for the Duplicator in the existential (j, k) -pebble game on (\mathbf{A}, \mathbf{B}) . Let $\mathcal{H} = \bigcup_{H_I \in \mathcal{M}} H_I$. Indeed, every mapping in \mathcal{H} is a projective homomorphism, and hence a partial homomorphism (with domain of size at most k). Moreover, the first property of a k -minimal family implies that \mathcal{H} is closed under restrictions and has the (j, k) -forth property.

(2) Assume that \mathbf{B} has relational width k . Take an arbitrary structure \mathbf{A} such that there exists a winning strategy \mathcal{H} for the Duplicator in the existential (k, k') -pebble game on (\mathbf{A}, \mathbf{B}) . We need to show that $\mathbf{A} \rightarrow \mathbf{B}$. By assumption, it suffices to exhibit a k -minimal family for (\mathbf{A}, \mathbf{B}) . For each $I \subseteq A$ with $|I| \leq k$, let $H_I = \{h \in \mathcal{H} \mid \text{dom}(h) = I\}$, and let \mathcal{M} be the family of all such sets H_I . From the properties of a winning strategy, it is easy to see that \mathcal{M} has property 1 of a k -minimal family. Now, for any $I \subseteq A$ with $|I| \leq k$, any $h \in H_I$, any $(n$ -ary) $R \in \tau$, and any $(a_1, \dots, a_n) \in R^{\mathbf{A}}$, the mapping $h|_{I \cap \{a_1, \dots, a_n\}}$ is a partial homomorphism which belongs to \mathcal{H} (by the restriction property of \mathcal{H}). Since $k' \geq \text{maxar}(\mathbf{B})$, it can be extended to a homomorphism $g : \mathbf{A}|_{\{a_1, \dots, a_n\}} \rightarrow \mathbf{B}$ (by the (k, k') -forth property of \mathcal{H}). Set $(b_1, \dots, b_n) = (g(a_1), \dots, g(a_n))$. It is easy to see that, since g and all its restrictions belong to \mathcal{H} , the obtained tuple (b_1, \dots, b_n) can be used to show that \mathcal{M} has property 2 of a k -minimal family.

Note that Theorem 19 shows that there is a correspondence between the parameters of relational width and treewidth duality, but it does not show how optimal parameters for these widths are related in general.

We will state results about bounded treewidth duality and bounded relational width in the way they were stated originally. By the above theorem, one can translate such results between the widths.

Similarly to Theorems 11 and 12, the (idempotent) polymorphisms determine whether a (core) structure has bounded treewidth duality.

Theorem 20 ([52]). *Let \mathbf{B}_1 and \mathbf{B}_2 be relational structures with the same universe and such that $\text{Pol}(\mathbf{B}_1) \subseteq \text{Pol}(\mathbf{B}_2)$. If $\text{co-CSP}(\mathbf{B}_2)$ is definable in Datalog, then so is $\text{co-CSP}(\mathbf{B}_1)$.*

Theorem 21 ([50]). *For a core structure \mathbf{B} , $\text{co-CSP}(\mathbf{B})$ is definable in Datalog if and only if $\text{co-CSP}(\mathbf{B}_c)$ is.*

We will now give examples of polymorphisms that guarantee that a structure has bounded treewidth duality.

Theorem 22 ([26], see also [41]). *If a structure \mathbf{B} has a $(l+1)$ -ary NU polymorphism then \mathbf{B} has l -treewidth duality.*

Tree duality is just a shorter name for 1-treewidth duality. It is known [26, 49] that every structure with tree duality has an obstruction set consisting of trees in the sense of Definition 11. (Note that this fact does not follow trivially from the definition of 1-treewidth duality.) In particular, if a structure \mathbf{B} has tree duality then it has $(1, \text{maxar}(\mathbf{B}))$ -treewidth duality. Structures with tree duality have been completely characterised in [26]. To state this result, we need to give a certain construction. For a τ -structure \mathbf{B} , its *power structure* is a τ -structure $\mathcal{P}_1(\mathbf{B})$ with universe consisting of all non-empty subsets of B , and, for each r -ary $R \in \tau$, we have $(A_1, \dots, A_r) \in R^{\mathcal{P}_1(\mathbf{B})}$ if and only if, for each $1 \leq i \leq r$ and each $a \in A_i$, there is $(a_1, \dots, a_r) \in R^{\mathbf{B}}$ such that $a_i = a$.

Theorem 23 ([26], see also [22]). *For any structure \mathbf{B} , the following conditions are equivalent:*

1. \mathbf{B} has tree duality.
2. The structure $\mathcal{P}_1(\mathbf{B})$ admits a homomorphism to \mathbf{B} .
3. For every $n \geq 2$, \mathbf{B} has an n -ary totally symmetric polymorphism.

Example 32. It is not difficult to see that, for $n \geq 2$, the operation $f_n = \bigwedge_{i=1}^n x_i$ is a TSI polymorphism of the structure \mathbf{B}_{kH} (see Example 27).

Theorem 19 implies that every structure with relational width 1 also has 1-treewidth duality. By using Theorem 23 and Definition 15, it is easy to show that, conversely, every structure with 1-treewidth duality has relational width 1. Note that, in general, the optimal parameters for relational width and treewidth duality need not be equal. For example, the structure $\mathbf{B} = \mathbf{K}_2$ of Example 9 has $(2,3)$ -treewidth duality, but not tree duality (since it has no binary TSI polymorphism). On the other hand, \mathbf{K}_2 has relational width 3 by Theorem 19, but not 2 (which can be seen by taking \mathbf{K}_3 as \mathbf{A}). In fact, it can be shown [20] that if a structure \mathbf{B} has relational width 2 then it has relational width 1.

Theorem 24 ([9]). *If a structure \mathbf{B} has a 2-semilattice polymorphism then \mathbf{B} has relational width 3.*

Some of the most studied varieties in universal algebra are the so-called *congruence distributive* varieties (see, e.g., [38, 42, 61]). For a core structure \mathbf{B} , the algebra $\mathbb{A}_{\mathbf{B}}$ belongs to a congruence distributive variety if, for some $n \geq 2$, \mathbf{B} has ternary polymorphisms d_0, \dots, d_n (called *Jónsson operations*) satisfying the identities from Theorem 16, and, in addition, such that $d_i(x, y, x) = x$ for all $0 \leq i \leq n$. In this case we say that $\mathbb{A}_{\mathbf{B}}$ is in the class $CD(n)$. Note that $\mathbb{A}_{\mathbf{B}}$ is in $CD(2)$ if and only if \mathbf{B} has a majority polymorphism (which is d_1 in this case).

Theorem 25 ([42]). *For any structure \mathbf{B} , if the algebra $\mathbb{A}_{\mathbf{B}}$ is in $CD(3)$ then \mathbf{B} has relational width $\min(|B|^2, \max(3, \text{maxar}(\mathbf{B})))$.*

Theorem 26 ([14]). *For any structure \mathbf{B} , if the algebra $\mathbb{A}_{\mathbf{B}}$ is in $CD(4)$ then \mathbf{B} has $(k - 1, k)$ -treewidth duality where $k = \max(3, \text{maxar}(\mathbf{B}))$.*

Theorem 20 makes it possible to introduce *algebras* having bounded treewidth duality: An algebra $\mathbb{A} = (B; F)$ has bounded treewidth duality if every structure \mathbf{B} with universe B such that $F \subseteq \text{Pol}(\mathbf{B})$ has bounded treewidth duality. The following result shows that bounded treewidth duality can be lifted further to varieties of algebras.

Theorem 27 ([52]). *If \mathbb{A} is an algebra with bounded treewidth duality then every finite algebra from the variety $\text{var}(\mathbb{A})$ also has bounded treewidth duality.*

Clearly, if $\text{CSP}(\mathbf{B})$ is **NP**-complete, then it does not have bounded treewidth duality unless **PTIME** = **NP**. Systems of linear equations (see Example 6 in this paper or the proof of Theorem 1 of [8]), as well as problems that can “simulate” them, provide benchmark examples of structures \mathbf{B} such that $\text{CSP}(\mathbf{B})$ is in **PTIME**, but \mathbf{B} does not have bounded treewidth duality [26]. Combining these two reasons for not having bounded treewidth duality, one obtains the following equivalent necessary conditions for bounded treewidth duality.

Theorem 28. *If a core structure \mathbf{B} has bounded treewidth duality then the following equivalent conditions hold:*

1. *The variety $\text{var}(\mathbb{A}_{\mathbf{B}})$ omits types 1 and 2.*
2. *There is $k \geq 2$ such that \mathbf{B} has n -ary weak NU polymorphisms for all $n \geq k$.*

In the above theorem, the necessity of condition (1) was proved in [52], and the equivalence of conditions (1) and (2) in [54]. It is shown in [51] that (the complement of) condition (1) is very closely related with the so-called property of “ability to count” which was introduced in [26] and conjectured there to be the main obstacle for a structure to have bounded treewidth duality.

Conjecture 1 ([52]). A core structure \mathbf{B} has bounded treewidth duality if and only if the equivalent conditions from Theorem 28 hold.

A somewhat different way of applying algebras to analyse a relational structure \mathbf{B} , via an *edge-coloured graph* $Gr(\mathbf{B})$ of the structure, was introduced in [8] (see also [13]). The conditions in Theorem 28 can be equivalently expressed in terms of properties of this graph, and a conjecture equivalent to Conjecture 1 was made in [8, 13].

Conjecture 1 was confirmed in the following important cases, and, interestingly, the best possible bound for some width turns out to be quite small.

Theorem 29 ([26]). *If \mathbf{B} is a 2-element core structure then \mathbf{B} has bounded treewidth if and only if \mathbf{B} has a semilattice polymorphism or a majority polymorphism. Moreover, in this case \mathbf{B} has 2-treewidth duality.*

A *factor* of an algebra \mathbb{A} is a homomorphic image of a subalgebra of \mathbb{A} .

Theorem 30 ([10]). *If \mathbf{B} is a core structure with $|B| \leq 3$ then \mathbf{B} has bounded relational width if and only if the algebra $\mathbb{A}_{\mathbf{B}}$ itself or each of its factors have an operation (depending on a factor) which is a majority operation or a 2-semilattice operation. Moreover, in this case \mathbf{B} has relational width 3.*

Theorem 31 ([7]). *Let \mathbf{B} be a structure containing all unary relations. Then \mathbf{B} has bounded relational width if and only if, for each two-element subset $C \subseteq B$, there is a polymorphism $f \in \text{Pol}(\mathbf{B})$ (depending on C) such that $f|_C$ is either a semilattice operation or a majority operation. Moreover, in this case \mathbf{B} has relational width 3.*

Conjecture 1 can be strengthened in the following sense. As we saw above, bounded treewidth duality is equivalent to expressibility in a certain infinitary logic. The expressive power of this logic is relatively weak, and it is natural to ask if it possible to express constraint satisfaction problems in terms of a more powerful logic. One such logic is $C_{\infty\omega}^\omega$ (see Section 2.3). This logic can express a number of undecidable problems (e.g., the HALTING problem). However, if Conjecture 1 is true than its expressive power for constraint satisfaction problems is no greater than that of Datalog.

Theorem 32 ([3]). *Let \mathbf{B} be a structure. If the variety $\text{var}(\mathbb{A}_{\mathbf{B}})$ admits type **1** or **2** then $\text{CSP}(\mathbf{B})$ is not expressible in $C_{\infty\omega}^\omega$.*

The complexity of checking the equivalent conditions from Theorem 28 is known. Therefore, under the assumption that Conjecture 1 is true, the complexity of checking if a structure or an algebra has bounded treewidth duality is also known.

Theorem 33 ([61]). *There is a polynomial time algorithm which, given a finite idempotent algebra \mathbb{A} , checks whether the variety $\text{var}(\mathbb{A})$ omits types **1** and **2**.*

Theorem 34 ([3]). *Assuming Conjecture 1 holds, the meta-problem for bounded treewidth duality is **NP**-hard for arbitrary structures, but tractable for structures of bounded size.*

Note that the tractability part of the above theorem can be derived from Theorem 33 following the strategy from [11] (where the case of omitting type **1** is treated).

It is a natural question to determine the complexity of recognising structures with j -treewidth duality for a fixed j . For $j = 1$, it follows from Theorem 23 that this problem is decidable, while the proof of Theorem 6.1 of [49] implies the following lower bound.

Theorem 35 ([49]). *It is **NP**-hard to decide whether a given structure \mathbf{B} has tree duality.*

In Section 4, we have considered the hierarchy problem for j -pathwidth dualities and found (see Theorem 17) that the hierarchy does not collapse. We now consider a similar problem for j -treewidth dualities. Let TW_j be the class of all structures with j -treewidth duality. Clearly, we have a hierarchy $TW_1 \subseteq TW_2 \subseteq TW_3 \subseteq TW_4 \subseteq \dots$.

It is easy to show that $TW_1 \subsetneq TW_2$. Consider the problem 2-COLOURABILITY, or $\text{CSP}(\mathbf{K}_2)$, from Example 9. The structure \mathbf{K}_2 has (2,3)-pathwidth duality, and hence (2,3)-treewidth duality. On the other hand, it is easy to see that \mathbf{K}_2 does not have a binary commutative polymorphism, and hence, by Theorem 23 (see also Example 15(1)), it cannot have 1-treewidth duality. Surprisingly, the question whether any other inclusion in the treewidth duality hierarchy is strict remains open. It may seem that Theorems 22, 25 and 26 contradict this claim. However, they give only an upper bound for the treewidth duality. For instance, every 2-element structure with an NU polymorphism has 2-treewidth duality.

6 Additional Remarks

6.1 Symmetric Datalog

A restriction of linear Datalog, *symmetric Datalog*, has been recently introduced in [23]. A linear Datalog program is called symmetric, if, for every rule of the form $t_0 : - t_1, t_2, \dots, t_n$ ($n \geq 1$), where t_0 and t_1 are IDBs, that appears in the program, the program also contains its “symmetric” rule $t_1 : - t_0, t_2, \dots, t_n$, obtained by formally swapping the IDBs in the rule. We say that $\text{co-CSP}(\mathbf{B})$ is definable in symmetric Datalog if it is accepted by a symmetric Datalog program. In broad terms, symmetric Datalog for CSP is to **LOGSPACE** what linear Datalog for CSP is to **NL**: if $\text{co-CSP}(\mathbf{B})$ is definable in symmetric Datalog then $\text{CSP}(\mathbf{B})$ is in **LOGSPACE**, and, for all problems $\text{CSP}(\mathbf{B})$ that are known to be in **LOGSPACE**, $\text{co-CSP}(\mathbf{B})$ is definable in symmetric Datalog [23]. In particular, this holds for all Boolean problems $\text{CSP}(\mathbf{B})$ in **LOGSPACE**. If \mathbf{B}_1 and \mathbf{B}_2 are structures such that $\text{Pol}(\mathbf{B}_1) \subseteq \text{Pol}(\mathbf{B}_2)$ and $\text{co-CSP}(\mathbf{B}_1)$ is definable in symmetric Datalog, then $\text{co-CSP}(\mathbf{B}_2)$ is also definable in symmetric Datalog (compare with Theorems 11 and 20). The following analog of Theorems 15 and 28 holds for symmetric Datalog: for a core structure \mathbf{B} , if $\text{co-CSP}(\mathbf{B})$ is definable in symmetric Datalog then the variety $\text{var}(\mathbf{A}_{\mathbf{B}})$ omits types **1**, **2**, **4**, and **5** (i.e., it admits only type **3**) [50]. It is shown in [23] that definability of $\text{co-CSP}(\mathbf{B})$ in symmetric Datalog is equivalent to definability in a certain fragment of second order logic (this parallels a result in [19]). It would be interesting to find a convenient pebble game and an appropriate notion of duality for symmetric Datalog, in the spirit of Theorems 10 and 18.

6.2 Extending Datalog with inequality and negation

One can extend (j, k) -Datalog, and the logic $\exists L_{\infty\omega}^{j,k}$ by allowing the use of inequalities (\neq) and negated atomic formulas (which must be EDBs in the case of

Datalog). The obtained logics are denoted (j, k) -Datalog(\neq, \neg) and $\exists L_{\infty\omega}^{j,k}(\neq, \neg)$, respectively. It was shown in [27] that these extensions do not add any expressive power for homomorphism-closed classes (e.g., such as $\text{co-CSP}(\mathbf{B})$). In other words, if a class $\text{co-CSP}(\mathbf{B})$ is definable in (j, k) -Datalog(\neq, \neg) then it is also definable in (j, k) -Datalog, and the same holds for $\exists L_{\infty\omega}^{j,k}(\neq, \neg)$. Moreover, a closer inspection of the proof reveals that this result remains true for linear (j, k) -Datalog and the logic $\exists M_{\infty\omega}^{j,k}$.

6.3 Infinite CSP

Up until now we have considered only finite structures. However, one can also consider the problem of deciding whether a given finite τ -structure admits a homomorphism to a fixed infinite τ -structure \mathbf{B} (see survey [4]). Some natural problems such as BETWEENNESS (see [29]) and the ACYCLICITY problem for digraphs can be represented as $\text{CSP}(\mathbf{B})$ for suitable infinitely countable structures \mathbf{B} (but not for any finite structure \mathbf{B}). Bounded treewidth duality for infinitely countable structures has been considered in [5, 6]. It was shown in these papers that, for general countable structures, Theorem 18 fails. However, there is a large class of structures (ω -categorical structures), for which Theorem 18 holds. Recall that a structure \mathbf{B} is called ω -categorical if, for each $n \geq 1$, there are only finitely many inequivalent first-order formulas with n free variables over \mathbf{B} . Moreover, analogs of Theorems 22 and 23 hold for such structures.

7 A list of open questions

1. If \mathbf{B} is a core structure with finite duality, how large can the minimal arity of its 1-tolerant NU polymorphism be?
2. Is the property of having j -pathwidth and j -treewidth duality for fixed j determined by the polymorphisms of a structure?
3. Is it true that a structure \mathbf{B} has bounded pathwidth duality whenever $\text{CSP}(\mathbf{B})$ is in **NL**?
4. Prove that every structure with an NU polymorphism has bounded pathwidth duality.
5. Are the conditions in Theorems 15 and 16 necessary and sufficient for a core structure to have bounded pathwidth duality?
6. Prove Conjecture 1 (that the conditions in Theorem 28 are necessary and sufficient for a core structure to have bounded treewidth duality).
7. For $j \geq 2$, is there a structure $\mathcal{P}_j(\mathbf{B})$ such that \mathbf{B} has j -treewidth duality if and only if $\mathcal{P}_j(\mathbf{B}) \rightarrow \mathbf{B}$ (particularly, for $j = 2$)?
8. Does the treewidth duality hierarchy collapse (in particular, to its second level) or not?
9. Are there structures that have bounded relational width, but not relational width 3?
10. Is it true that the number k from Theorem 28 can always be chosen to be equal to 3?

11. Prove that a structure \mathbf{B} has bounded treewidth (or even bounded path-width) duality whenever the algebra $\mathbb{A}_{\mathbf{B}}$ is in $CD(n)$ for some n ?
12. Find a pebble-game and a duality characterisation for structures \mathbf{B} such that $\text{co-CSP}(\mathbf{B})$ is definable in symmetric Datalog.
13. Is it true that $\text{co-CSP}(\mathbf{B})$ is definable in symmetric Datalog whenever $\text{CSP}(\mathbf{B})$ is in **LOGSPACE**?
14. Is it true that, for a core structure \mathbf{B} , $\text{co-CSP}(\mathbf{B})$ is definable in symmetric Datalog whenever the variety $\text{var}(\mathbb{A}_{\mathbf{B}})$ admits only type **3**?
15. Are there other naturally arising dualities for the CSP?

8 Acknowledgements

The authors would like to thank Catarina Carvalho, Víctor Dalmau, Pavol Hell, Jarik Nešetřil, Claude Tardif and Pascal Tesson for useful discussions and helpful comments. Remarks from an anonymous reviewer are also appreciated.

References

1. F. Afrati and S. Cosmodakis. Expressiveness of restricted recursive queries. In *STOC'89*, pages 113–126, 1989.
2. A. Atserias. On digraph coloring problems and treewidth duality. In *LICS'05*, pages 106–115, 2005.
3. A. Atserias, A. Bulatov, and A. Dawar. Affine systems of equations and counting infinitary logic. In *ICALP'07*, volume 4596 of *LNCS*, pages 558–570, 2007.
4. M. Bodirsky. Constraint satisfaction problems with infinite templates. In *The Same Volume*. 2008.
5. M. Bodirsky and V. Dalmau. Datalog and constraint satisfaction with infinite templates. In *STACS'06*, volume 3884 of *LNCS*, pages 646–659, 2006.
6. M. Bodirsky and V. Dalmau. Datalog and constraint satisfaction with infinite templates, 2007. manuscript, full version.
7. A. Bulatov. Tractable conservative constraint satisfaction problems. In *LICS'03*, pages 321–330, 2003.
8. A. Bulatov. A graph of a relational structure and constraint satisfaction problems. In *LICS'04*, pages 448–457, 2004.
9. A. Bulatov. Combinatorial problems raised from 2-semilattices. *Journal of Algebra*, 298(2):321–339, 2006.
10. A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *Journal of the ACM*, 53(1):66–120, 2006.
11. A. Bulatov and P. Jeavons. Algebraic structures in combinatorial problems. Technical Report MATH-AL-4-2001, Technische Universität Dresden, Germany, 2001.
12. A. Bulatov, P. Jeavons, and A. Krokhin. Classifying complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34(3):720–742, 2005.
13. A. Bulatov and M. Valeriote. Recent results on the algebraic approach to the CSP. In *The Same Volume*. 2008.
14. C. Carvalho, V. Dalmau, P. Marković, and M. Maróti. $CD(4)$ has bounded width. *Algebra Universalis*. accepted.
15. M. Clasen and M. Valeriote. Tame congruence theory. In *Lectures on Algebraic Model Theory*, volume 15 of *Fields Institute Monographs*, pages 67–111. 2002.

16. D. Cohen and P. Jeavons. The complexity of constraint languages. In *Handbook of Constraint Programming*, chapter 8. Elsevier, 2006.
17. M.C. Cooper, D.A. Cohen, and P.G. Jeavons. Characterising tractable constraints. *Artificial Intelligence*, 65:347–361, 1994.
18. 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.
19. V. Dalmau. Linear Datalog and bounded path duality for relational structures. *Logical Methods in Computer Science*, 1(1), 2005. (electronic).
20. V. Dalmau. A note on relational width. manuscript, 2007.
21. V. Dalmau and A. Krokhin. Majority constraints have bounded pathwidth duality. *European Journal of Combinatorics*, 2007. (accepted).
22. V. Dalmau and J. Pearson. Set functions and width 1 problems. In *CP'99*, volume 1713 of *LNCS*, pages 159–173. Springer-Verlag, 1999.
23. L. Egri, B. Larose, and P. Tesson. Symmetric Datalog and constraint satisfaction problems in Logspace. In *LICS'07*, pages 193–202, 2007.
24. T. Feder. Classification of homomorphisms to oriented cycles and of k -partite satisfiability. *SIAM Journal on Discrete Mathematics*, 14(4):471–480, 2001.
25. T. Feder, P. Hell, and J. Huang. Bi-arc graphs and the complexity of list homomorphisms. *Journal of Graph Theory*, 42:61–80, 2003.
26. 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.
27. T. Feder and M.Y. Vardi. Homomorphism closed vs. existential positive. In *LICS'03*, pages 311–320, 2003.
28. M. Furst, J. Saxe, and M. Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, 1984.
29. M. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, CA., 1979.
30. P. Hell. From graph colouring to constraint satisfaction: there and back again. In *Topics in Discrete Mathematics*, volume 26 of *Algorithms and Combinatorics*, pages 407–432. Springer, 2006.
31. P. Hell and J. Nešetřil. On the complexity of H -coloring. *Journal of Combinatorial Theory, Ser.B*, 48:92–110, 1990.
32. P. Hell and J. Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.
33. P. Hell, J. Nešetřil, and X. Zhu. Duality and polynomial testing of tree homomorphisms. *Trans. Amer. Math. Soc.*, 348:147–156, 1996.
34. P. Hell, J. Nešetřil, and X. Zhu. Duality of graph homomorphisms. In *Combinatorics, Paul Erdős is Eighty (Vol.2)*, volume 2 of *Bolyai Soc. Math. Stud.*, pages 271–282. János Bolyai Math. Soc., 1996.
35. P. Hell, H. Zhou, and X. Zhu. On homomorphisms to acyclic local tournaments. *Journal of Graph Theory*, 20(4):467–471, 1995.
36. P. Hell and X. Zhu. Homomorphisms to oriented paths. *Discrete Mathematics*, 132:107–114, 1994.
37. P. Hell and X. Zhu. The existence of homomorphisms to oriented cycles. *SIAM Journal on Discrete Mathematics*, 8:208–222, 1995.
38. D. Hobby and R.N. McKenzie. *The Structure of Finite Algebras*, volume 76 of *Contemporary Mathematics*. AMS, Providence, R.I., 1988.

39. P. Idziak, P. Markovic, R. McKenzie, M. Valeriote, and R. Willard. Tractability and learnability arising from algebras with few subpowers. In *LICS'07*, pages 213–222, 2007.
40. P. Jeavons. On the algebraic structure of combinatorial problems. *Theoretical Computer Science*, 200:185–204, 1998.
41. P.G. Jeavons, D.A. Cohen, and M.C. Cooper. Constraints, consistency and closure. *Artificial Intelligence*, 101(1-2):251–265, 1998.
42. E.W. Kiss and M. Valeriote. On tractability and congruence distributivity. *Logical Methods in Computer Science*, 3(2), 2007. (electronic).
43. Ph.G. Kolaitis. On the expressive power of logics on finite models. In *Finite Model Theory and its Applications*. Springer, 2007.
44. Ph.G. Kolaitis and M.Y. Vardi. On the expressive power of Datalog: tools and a case study. *Journal of Computer and System Sciences*, 51:110–134, 1995.
45. Ph.G. Kolaitis and M.Y. Vardi. Conjunctive-query containment and constraint satisfaction. *Journal of Computer and System Sciences*, 61:302–332, 2000.
46. Ph.G. Kolaitis and M.Y. Vardi. A logical approach to constraint satisfaction. In *Finite Model Theory and its Applications*. Springer, 2007.
47. P. Komárek. Some new good characterisations of directed graphs. *Časopis Pěst. Mat.*, 51:348–354, 1984.
48. A. Krokhin, A. Bulatov, and P. Jeavons. The complexity of constraint satisfaction: an algebraic approach. In *Structural Theory of Automata, Semigroups, and Universal Algebra*, pages 181–213. Springer Verlag, 2005.
49. B. Larose, C. Loten, and C. Tardif. A characterisation of first-order constraint satisfaction problems. *Logical Methods in Computer Science*, 3(4), 2007. (electronic).
50. B. Larose and P. Tesson. Universal algebra and hardness results for constraint satisfaction problems. In *ICALP'07*, volume 4596 of *LNCS*, pages 267–278, 2007.
51. B. Larose, M. Valeriote, and L. Zádori. An algebraic characterisation of the ability to count. manuscript, 2008.
52. B. Larose and L. Zádori. Bounded width problems and algebras. *Algebra Universalis*, 56(3-4):439–466, 2007.
53. L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
54. M. Maróti and R. McKenzie. Existence theorems for weakly symmetric operations. *Algebra Universalis*, 2007. to appear.
55. J. Nešetřil and A. Pultr. On classes of relations and graphs determined by subobjects and factorobjects. *Discrete Mathematics*, 22:287–300, 1978.
56. 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.
57. J. Nešetřil and C. Tardif. Short answers to exponentially long questions: extremal aspects of homomorphism duality. *SIAM Journal on Discrete Mathematics*, 19(4):914–920, 2005.
58. F. Rossi, P. van Beek, and T. Walsh, editors. *Handbook of Constraint Programming*. Elsevier, 2006.
59. B. Rossman. Existential positive types and preservation under homomorphisms. In *LICS'05*, pages 467–476, 2005.
60. A. Szendrei. *Clones in Universal Algebra*, volume 99 of *Seminaires de Mathématiques Supérieures*. University of Montreal, 1986.
61. M. Valeriote. A subalgebra intersection property for congruence-distributive varieties. *Canadian Journal of Mathematics*, 2007. to appear.
62. L. Zádori. Posets, near-unanimity functions and zigzags. *Bulletin of the Australian Mathematical Society*, 47:79–93, 1993.