

Enumerating Homomorphisms[☆]

Andrei A. Bulatov

School of Computing Science, Simon Fraser University, Burnaby, Canada

Víctor Dalmau^{1,*}

Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

Martin Grohe

Institut für Informatik, Humboldt-Universität, Berlin, Germany

Dániel Marx²

Institut für Informatik, Humboldt-Universität, Berlin, Germany

Abstract

The homomorphism problem for relational structures is an abstract way of formulating constraint satisfaction problems (CSP) and various problems in database theory. The decision version of the homomorphism problem received a lot of attention in literature; in particular, the way the graph-theoretical structure of the variables and constraints influences the complexity of the problem is intensively studied. Here we study the problem of enumerating all the solutions with polynomial delay from a similar point of view. It turns out that the enumeration problem behaves very differently from the decision version. We give evidence that it is unlikely that a characterization result similar to the decision version can be obtained. Nevertheless, we show nontrivial cases where enumeration can be done with polynomial delay.

Keywords: enumeration, computational complexity, homomorphisms, constraint satisfaction

[☆]A preliminary version of this paper was published in the proceedings of STACS'09

*Corresponding author.

Email addresses: abulatov@cs.sfu.ca (Andrei A. Bulatov), victor.dalmau@upf.edu (Víctor Dalmau), grohe@informatik.hu-berlin.de (Martin Grohe), dmarx@cs.bme.hu (Dániel Marx)

¹Supported by the MCyT through grants TIN2006-15387-C03-03 and TIN 2004-04343, and the program José Castillejo.

²Supported by the Alexander von Humboldt Foundation.

1. Introduction

Constraint satisfaction problems (CSP) form a rich class of algorithmic problems with applications in many areas of computer science. We only mention database systems, where CSPs appear in the guise of the conjunctive query containment problem and the closely related problem of evaluating conjunctive queries. It has been observed by Feder and Vardi [13] that as abstract problems, CSPs are homomorphism problems for relational structures. Algorithms for and the complexity of constraint satisfaction problems have been intensely studied (e.g. [18, 9, 3, 4]), not only for the standard decision problems but also optimization versions (e.g. [2, 21, 22, 24]) and counting versions (e.g. [10, 6, 7, 12]) of CSPs.

In this paper we study the *CSP enumeration problem*, that is, problem of computing all solutions for a given CSP instance. More specifically, we are interested in the question which structural restrictions on CSP instances guarantee tractable enumeration problems. “Structural restrictions” are restrictions on the structure induced by the constraints on the variables. Examples of structural restrictions are “every variable occurs in at most 5 constraints” and “the constraints form an acyclic hypergraph³”. This can most easily be made precise if we view CSPs as homomorphism problems: Given two relational structures \mathbb{A}, \mathbb{B} , decide if there is a homomorphism from \mathbb{A} to \mathbb{B} . Here the elements of the structure \mathbb{A} correspond to the variables of the CSP and the elements of the structure \mathbb{B} correspond to the possible values. Structural restrictions are restrictions on the structure \mathbb{A} . If \mathcal{A} is a class of structures, then $\text{CSP}(\mathcal{A}, -)$ denotes the restriction of the general CSP (or homomorphism problem) where the “left hand side” input structure \mathbb{A} is taken from the class \mathcal{A} . $\text{ECSP}(\mathcal{A}, -)$ denotes the corresponding enumeration problem: Given two relational structures $\mathbb{A} \in \mathcal{A}$ and \mathbb{B} , compute the set of all homomorphisms from \mathbb{A} to \mathbb{B} . The enumeration problem is of particular interest in the database context, where we are usually not only interested in the question of whether the answer to a query is nonempty, but want to compute all tuples in the answer. We will also briefly discuss the corresponding *search* problem, denoted $\text{SCSP}(\mathcal{A}, -)$: Find a solution if one exists.

It has been shown in [1] that $\text{ECSP}(\mathcal{A}, -)$ can be solved in polynomial time if and only if the number of solutions (that is, homomorphisms) for all instances is polynomially bounded in terms of the input size and that this is the case if and only if the structures in the class \mathcal{A} have bounded fractional edge cover number. However, usually we cannot expect the number of solutions to be polynomial. In this case, we may ask which conditions on \mathcal{A} guarantee that $\text{ECSP}(\mathcal{A}, -)$ has a polynomial delay algorithm. A *polynomial delay algorithm* for an enumeration problem is required to produce the first solution in polynomial time and then iteratively compute all solutions (each solution only once), leaving only polynomial time between two successive solutions. In particular, this guarantees that the algorithm computes all solutions in *polynomial total time*, that is, in time polynomial in the input size plus output size.

³The other type of restrictions studied in the literature on CSP are “constraint language restrictions”, that is, restrictions on the structure imposed by the constraint relations on the values. An example of a constraint language restriction is “all clauses of a SAT instance, viewed as a Boolean CSP, are Horn clauses”.

It is easy to see that $\text{ECSP}(\mathcal{A}, -)$ has a polynomial delay algorithm if the class \mathcal{A} has bounded tree width. It is also easy to see that there are classes \mathcal{A} of unbounded tree width such that $\text{ECSP}(\mathcal{A}, -)$ has a polynomial delay algorithm. It follows from our results that examples of such classes are the class of all grids or the class of all complete graphs with a loop on every vertex. It is known that the decision problem $\text{CSP}(\mathcal{A}, -)$ is in polynomial time if and only if the cores of the structures in \mathcal{A} have bounded tree width [16] (provided the arity of the constraints is bounded, and under some reasonable complexity theoretic assumptions). A *core* of a relational structure \mathbb{A} is a minimal substructure $\mathbb{A}' \subseteq \mathbb{A}$ such that there is a homomorphism from \mathbb{A} to \mathbb{A}' ; minimality is with respect to inclusion. It is easy to see that all cores of a structure are isomorphic. Hence we usually speak of “the” core of a structure. Note that the core of a grid (and of any other bipartite graph with at least one edge) is a single edge, and the core of a complete graph with all loops present (and of any other graph with a loop) is a single vertex with a loop on it. The core of a complete graph with no loops is the graph itself. As a polynomial delay algorithm for an enumeration algorithms yields a polynomial time algorithm for the corresponding decision problem, it follows that $\text{ECSP}(\mathcal{A}, -)$ can only have a polynomial delay algorithm if the cores of the structures in \mathcal{A} have bounded tree width. Unfortunately, there are examples of classes \mathcal{A} that have cores of bounded tree width, but for which $\text{ECSP}(\mathcal{A}, -)$ has no polynomial delay algorithm unless $P = NP$ (see Example 1).

Our main algorithmic results show that $\text{ECSP}(\mathcal{A}, -)$ has a polynomial delay algorithm if the cores of the structures in \mathcal{A} have bounded tree width and if, in addition, they can be reached in a sequence of “small steps.” An *endomorphism* of a structure is a homomorphism of a structure to itself. A *retraction* is an endomorphism that is the identity mapping on its image. Every structure has a retraction to its core. However, in general, the only way to map a structure to its core may be by collapsing the whole structure at once. As an example, consider a path with a loop on both endpoints. The core consists of a single vertex with a loop. (More precisely, the two cores are the two endpoints with their loops.) The only endomorphism of this structure to a proper substructure maps the whole structure to its core. Compare this with a path that only has a loop on one endpoint. Again, the core is a single vertex with a loop, but now we can reach the core by a sequence of retractions, mapping a path of length n to a subpath of length $n - 1$ and then to a subpath of length $n - 2$ et cetera. We prove that if \mathcal{A} is a class of structures whose cores have bounded tree width and can be reached by a sequence of retractions each of which only moves a bounded number of vertices, then $\text{ECSP}(\mathcal{A}, -)$ has a polynomial delay algorithm (Theorem 3).

We also consider more general sequences of retractions or endomorphism from a structure to its core. We say that a sequence of endomorphisms from a structure \mathbb{A}_0 to a substructure $\mathbb{A}_1 \subset \mathbb{A}_0$, from \mathbb{A}_1 to a substructure \mathbb{A}_2 , \dots , to a structure \mathbb{A}_n has *bounded width* if \mathbb{A}_n and, for each $i \leq n$, the “difference between \mathbb{A}_i and \mathbb{A}_{i-1} ” has bounded tree width. We prove that if we are given a sequence of endomorphisms of bounded width together with the input structure \mathbb{A} , then we can compute all solutions by a polynomial delay algorithm. Unfortunately, in general we cannot compute such a sequence of endomorphisms efficiently. We prove that even for width 1 it is NP-complete to decide whether such a sequence exists (Theorem 2). We also show that the existence of a sequence of bounded width endomorphisms is not a sufficient criterion

for tractability if this sequence is not explicitly given (Proposition 1).

In the last section, we briefly discuss the problem of computing projections of solutions of a CSP, which is equivalent to the problem of evaluating conjunctive queries in relational databases. This problem is significantly harder, and we show that our criteria for tractability beyond bounded tree width cannot be extended (Example 3).

Finally, we remark that our results are far from giving a complete classification of the classes \mathcal{A} for which $\text{ECSP}(\mathcal{A}, -)$ has a polynomial delay algorithm and those classes for which it does not. Indeed, we show that it will be difficult to obtain such a classification, because such a classification would imply a solution to the notoriously open *CSP dichotomy conjecture* of Feder and Vardi [13] (see Section 3 for details).

2. Preliminaries

Relational structures. A *vocabulary* τ is a finite set of *relation symbols* of specified arities. A *relational structure* \mathbb{A} over τ consists of a finite set A called the *universe* of \mathbb{A} and for each relation symbol $R \in \tau$, say, of arity r , an r -ary relation $R^{\mathbb{A}} \subseteq A^r$. Note that we require vocabularies and structures to be finite. A structure \mathbb{A} is a *substructure* of a structure \mathbb{B} if $A \subseteq B$ and $R^{\mathbb{A}} \subseteq R^{\mathbb{B}}$ for all $R \in \tau$. We write $\mathbb{A} \subseteq \mathbb{B}$ to denote that \mathbb{A} is a substructure of \mathbb{B} and $\mathbb{A} \subset \mathbb{B}$ to denote that \mathbb{A} is a *proper* substructure of \mathbb{B} , that is, $\mathbb{A} \subseteq \mathbb{B}$ and $\mathbb{A} \neq \mathbb{B}$. A substructure $\mathbb{A} \subseteq \mathbb{B}$ is *induced* if for all $R \in \tau$, say, of arity r , we have $R^{\mathbb{A}} = R^{\mathbb{B}} \cap A^r$. For a subset $A \subseteq B$, we write $\mathbb{B}[A]$ to denote the induced substructure of \mathbb{B} with universe A .

Homomorphisms. We often abbreviate tuples (a_1, \dots, a_k) by \mathbf{a} . If φ is a mapping whose domain contains a_1, \dots, a_k we write $\varphi(\mathbf{a})$ to abbreviate $(\varphi(a_1), \dots, \varphi(a_k))$. A *homomorphism* from a relational structure \mathbb{A} to a relational structure \mathbb{B} is a mapping $\varphi : A \rightarrow B$ such that for all $R \in \tau$ and all tuples $\mathbf{a} \in R^{\mathbb{A}}$ we have $\varphi(\mathbf{a}) \in R^{\mathbb{B}}$. An *endomorphism* of \mathbb{A} is a homomorphism from \mathbb{A} to \mathbb{A} , and a *retraction* of \mathbb{A} is an endomorphism h such that for all $a \in A$ it holds that $h(h(a)) = h(a)$. A *partial homomorphism* on $C \subseteq A$ to \mathbb{B} is a homomorphism of $\mathbb{A}[C]$ to \mathbb{B} . It is sometimes useful when designing examples to exclude certain homomorphisms or endomorphisms. The simplest way to do that is to use unary relations. For example, if R is a unary relation and $(a) \in R^{\mathbb{A}}$ we say that a has color R . Now if $b \in B$ does not have color R then no homomorphism from \mathbb{A} to \mathbb{B} maps a to b . If \mathbb{A} is a τ -structure and φ is a mapping with $\text{dom}(\varphi) = A$, then $\varphi(\mathbb{A})$ is the τ -structure with universe $\varphi(A)$ and with $R^{\varphi(\mathbb{A})} = \{\varphi(\mathbf{a}) \mid \mathbf{a} \in R^{\mathbb{A}}\}$. Note that a mapping $\varphi : A \rightarrow B$ is a homomorphism from \mathbb{A} to \mathbb{B} if and only if $\varphi(\mathbb{A})$ is a substructure (not necessarily induced) of \mathbb{B} .

Two structures \mathbb{A} and \mathbb{B} are *homomorphically equivalent* if there is a homomorphism from \mathbb{A} to \mathbb{B} and also a homomorphism from \mathbb{B} to \mathbb{A} . Note that if structures \mathbb{A} and \mathbb{A}' are homomorphically equivalent, then for every structure \mathbb{B} there is a homomorphism from \mathbb{A} to \mathbb{B} if and only if there is a homomorphism from \mathbb{A}' to \mathbb{B} ; in other words: the instances (\mathbb{A}, \mathbb{B}) and $(\mathbb{A}', \mathbb{B})$ of the decision CSP are equivalent. However, the two instances may have vastly different sizes, and the complexity of solving the search and enumeration problems for them can also be quite different. Homomorphic equivalence is closely related to the concept of the core of a structure: A structure \mathbb{A} is a *core* if there is no homomorphism from \mathbb{A} to a proper substructure of \mathbb{A} . A core

of a structure \mathbb{A} is a substructure $\mathbb{A}' \subseteq \mathbb{A}$ such that there is a homomorphism from \mathbb{A} to \mathbb{A}' and \mathbb{A}' is a core. Equivalently, a core of a structure can be defined as a minimal substructure \mathbb{A}' of \mathbb{A} such that there is a homomorphism from \mathbb{A} to \mathbb{A}' . Obviously, every core of a structure is homomorphically equivalent to the structure. We observe another basic fact about cores:

Observation 1. Let \mathbb{A} and \mathbb{B} be homomorphically equivalent structures, and let \mathbb{A}' and \mathbb{B}' be cores of \mathbb{A} and \mathbb{B} , respectively. Then \mathbb{A}' and \mathbb{B}' are isomorphic. In particular, all cores of a structure \mathbb{A} are isomorphic. Therefore, we often speak of *the* core of \mathbb{A} .

Observation 2. It is easy to see that it is NP-hard to decide, given structures $\mathbb{A} \subseteq \mathbb{B}$, whether \mathbb{A} is isomorphic to the core of \mathbb{B} . (For an arbitrary graph G , let \mathbb{A} be a triangle and \mathbb{B} the disjoint union of G with \mathbb{A} . Then \mathbb{A} is a core of \mathbb{B} if and only if G is 3-colorable.) Hell and Nešetřil [17] proved that it is co-NP-complete to decide whether a graph is a core.

Tree decompositions. A *tree decomposition* of a graph G is a pair (T, B) , where T is a tree and B is a mapping that associates with every node $t \in V(T)$ a set $B_t \subseteq V(G)$ such that (1) for every $v \in V(G)$ the set $\{t \in V(T) \mid v \in B_t\}$ is connected in T , and (2) for every $e \in E(G)$ there is a $t \in V(T)$ such that $e \subseteq B_t$. The sets B_t , for $t \in V(T)$, are called the *bags* of the decomposition. It is sometimes convenient to have the tree T in a tree decomposition rooted; we always assume it is. The *width* of a tree decomposition (T, B) is $\max\{|B_t| \mid t \in V(T)\} - 1$. The *tree width* of a graph G , denoted by $\text{tw}(G)$, is the minimum of the widths of all tree decompositions of G .

We need to transfer some of the notions of graph theory to arbitrary relational structures. The *Gaifman graph* (also known as *primal graph*) of a relational structure \mathbb{A} over vocabulary τ is the graph $G(\mathbb{A})$ with vertex set A and an edge between a and b if $a \neq b$ and there is a relation symbol $R \in \tau$, say, of arity r , and a tuple $(a_1, \dots, a_r) \in R^{\mathbb{A}}$ such that $a, b \in \{a_1, \dots, a_r\}$. We can now transfer graph-theoretic notions to relational structures. In particular, a subset $B \subseteq A$ is *connected* in a structure \mathbb{A} if it is connected in $G(\mathbb{A})$. A *tree decomposition* of a structure \mathbb{A} can simply be defined to be a tree-decomposition of $G(\mathbb{A})$. Equivalently, a tree decomposition of \mathbb{A} can be defined directly by replacing the second condition in the definition of tree decompositions of graphs by (2') for every $R \in \tau$ and $(a_1, \dots, a_r) \in R^{\mathbb{A}}$ there is a $t \in V(T)$ such that $\{a_1, \dots, a_r\} \subseteq B_t$. A class \mathcal{C} of structures has *bounded tree width* if there is a $w \in \mathbb{N}$ such that $\text{tw}(\mathbb{A}) \leq w$ for all $\mathbb{A} \in \mathcal{C}$. A class \mathcal{C} of structures has *bounded tree width modulo homomorphic equivalence* if there is a $w \in \mathbb{N}$ such that every $\mathbb{A} \in \mathcal{C}$ is homomorphically equivalent to a structure of tree width at most w .

Observation 3. A structure \mathbb{A} is homomorphically equivalent to a structure of tree width at most w if and only if the core of \mathbb{A} has tree width at most w .

The Constraint Satisfaction Problem. For two classes \mathcal{A} and \mathcal{B} of structures, the *Constraint Satisfaction Problem*, $\text{CSP}(\mathcal{A}, \mathcal{B})$, is the following problem:

$\text{CSP}(\mathcal{A}, \mathcal{B})$
Instance: $\mathbb{A} \in \mathcal{A}, \mathbb{B} \in \mathcal{B}$
Problem: Decide if there is a homomorphism from \mathbb{A} to \mathbb{B} .

The CSP is a decision problem. The variation of it we study in this paper is the following enumeration problem:

ECSP(\mathcal{A}, \mathcal{B})
Instance: $\mathbb{A} \in \mathcal{A}, \mathbb{B} \in \mathcal{B}$
Problem: Output all the homomorphisms from \mathbb{A} to \mathbb{B} .

We shall also refer to the search problem, SCSP(\mathcal{A}, \mathcal{B}), in which the goal is to find one solution of a CSP-instance or output ‘no’ if a solution does not exist.

If one of the classes \mathcal{A}, \mathcal{B} is the class of all finite structures, then we denote the corresponding CSPs by CSP($\mathcal{A}, -$), CSP($- , \mathcal{B}$) (respectively, ECSP($\mathcal{A}, -$), ECSP($- , \mathcal{B}$), SCSP($\mathcal{A}, -$), SCSP($- , \mathcal{B}$)).

The decision CSP has been intensely studied. In particular it has been shown, under standard complexity-theoretic assumptions, that if a class \mathcal{C} of structures has bounded arity then CSP($\mathcal{C}, -$) is solvable in polynomial time if and only if \mathcal{C} has bounded tree width modulo homomorphic equivalence [16] whereas if the arity of \mathcal{C} is not bounded then CSP($\mathcal{C}, -$) is fixed-parameter tractable if and only if \mathcal{C} has bounded submodular width [23].

Problems of the form CSP($- , \mathcal{C}$) have been studied mostly in the case when \mathcal{C} is 1-element. Problems of this type are sometimes referred to as *non-uniform*. It is conjectured that every non-uniform problem is either solvable in polynomial time or NP-complete (the so-called *Dichotomy Conjecture*) [13]. Although this conjecture is proved in several particular cases [18, 8, 9, 3], in its general form it is believed to be very difficult.

A search CSP is clearly no easier than the corresponding decision problem. While any non-uniform search problem SCSP($- , \mathcal{C}$) is polynomial time reducible to its decision version CSP($- , \mathcal{C}$) [11], nothing is known about the complexity of search problems SCSP($\mathcal{C}, -$) except the result we state in Section 3. Regarding enumeration of CSPs, some initial results on the complexity of non-uniform enumerating problems have been reported in [26]. Also, the question of enumerating solutions “projected” over a given set of variables has been investigated in [15].

3. Tractable structures for enumeration

Since even an easy CSP may have exponentially many solutions, the model of choice for ‘easy’ enumeration problems is algorithms with polynomial delay [20]. An algorithm Alg is said to solve a CSP *with polynomial delay* (w.p.d. for short) if there is a polynomial $p(n)$ such that, for every instance of size n , Alg outputs ‘no’ in a time bounded by $p(n)$ if there is no solution, otherwise it generates all solutions to the instance such that no solution is output twice, the first solution is output after at most $p(n)$ steps after the computation starts, and time between outputting two consequent solutions does not exceed $p(n)$.

If a class of relational structures \mathcal{C} has bounded arity, the aforementioned result of Grohe [16] imposes strong restrictions on enumeration problems solvable w.p.d.

Observation 4. If a class of relational structures \mathcal{C} with bounded arity does not have bounded tree width modulo homomorphic equivalence, then $\text{ECSP}(\mathcal{C}, -)$ is not solvable w.p.d., unless $\text{FPT} = W[1]$.

Unlike for the decision version, the converse is not true. Indeed, the following example shows that bounded tree width modulo homomorphic equivalence does not imply enumerability w.p.d. This also has been noted in [25].

Example 1. Let \mathbb{A}_k be the disjoint union of a k -clique and a loop and let $\mathcal{A} = \{\mathbb{A}_k \mid k \geq 1\}$. Clearly, the core of each graph in \mathcal{A} has bounded tree width (in fact, it is a single element), hence $\text{CSP}(\mathcal{A}, -)$ is polynomial-time solvable. For an arbitrary graph \mathbb{B} without loops, let \mathbb{B}' be the disjoint union of \mathbb{B} and a loop. It is clear that there is always a trivial homomorphism from \mathbb{A}_k (for any $k \geq 1$) to \mathbb{B}' that maps everything into the loop. There exist homomorphisms different from the trivial one if and only if \mathbb{B} contains a k -clique. Thus if we are able to check in polynomial time whether there is a second homomorphism, then we are able to test if \mathbb{B} has a k -clique. Therefore, although $\text{CSP}(\mathcal{A}, -)$ and $\text{SCSP}(\mathcal{A}, -)$ are polynomial-time solvable, a w.p.d. enumeration algorithm for $\text{ECSP}(\mathcal{A}, -)$ would imply $\text{P} = \text{NP}$.

It is not difficult to show that $\text{ECSP}(\mathcal{C}, -)$ is enumerable w.p.d. if \mathcal{C} has bounded tree width. Instead of giving a direct proof we shall derive this from a more general result in Section 6.

Thus enumerability w.p.d. has a different tractability criterion than the decision version, and this criterion lies somewhere between bounded tree width and bounded tree width modulo homomorphic equivalence. Thus in order to ensure that the solutions can be enumerated w.p.d., we have to make further restrictions on the way the structure can be mapped to its bounded tree width core. The main new definition of the paper requires that the core is reached by “small steps”:

Definition 1. Let \mathbb{A} be a relational structure with universe A . We say that \mathbb{A} has a sequence of endomorphisms of width k if there are subsets $A = A_0 \supset A_1 \supset \dots \supset A_n \neq \emptyset$ and homomorphisms $\varphi_1, \dots, \varphi_n$ such that

1. φ_i is a homomorphism from $\mathbb{A}[A_{i-1}]$ to $\mathbb{A}[A_i]$,
2. $\varphi_i(A_{i-1}) = A_i$ for $1 \leq i \leq n$;
3. for every $0 \leq i < n$, the tree width of $\mathbb{A}[A_i \setminus A_{i+1}]$ is at most k ;
4. $\mathbb{A}[A_n]$ has tree width at most k .

If $\varphi_1, \dots, \varphi_n$ are retractions of $\mathbb{A}[A_0], \dots, \mathbb{A}[A_{n-1}]$, then we say that \mathbb{A} has a sequence of retractions of width k .

In Section 4, we show that enumeration for (\mathbb{A}, \mathbb{B}) can be done w.p.d. if a sequence of bounded width endomorphisms for \mathbb{A} is given in the input. Unfortunately, we cannot claim that $\text{ECSP}(\mathcal{A}, -)$ can be done w.p.d. if every structure in \mathcal{A} has such a sequence, since we do not know how to find such sequences efficiently. In fact, as we show in Section 5, it is hard to check if a width-1 sequence exists for a given structure. Furthermore, we construct a class \mathcal{A} where every structure has a width-2 sequence, but $\text{ECSP}(\mathcal{A}, -)$ cannot be done w.p.d., unless $\text{P} = \text{NP}$. This means that it is not possible

to get around the problem of not being able to find the sequences (for example, by finding sequences with somewhat larger width or by constructing the sequence during the enumeration). Thus having a bounded width sequence of endomorphisms is not the right tractability criterion. We then investigate a more restrictive notion, where the bound is not on the tree width of the difference of the layers but on the number of elements in the differences and show that this yields enumeration w.p.d.

However, in the rest of the section, we give evidence that enumeration problems solvable w.p.d. cannot be characterized in simple terms. Indeed, a description of enumeration problems solvable w.p.d. would imply a description of non-uniform decision problems solvable in polynomial time. This is shown via an analogous result for the search version of the problem, which might be of independent interest. By $\mathbb{A} \oplus \mathbb{B}$ we denote the disjoint union of relational structures \mathbb{A} and \mathbb{B} .

Lemma 1. *Let \mathbb{B} be a relational structure that is a core, and let $\mathcal{C}_{\mathbb{B}}$ be $\{\mathbb{A} \oplus \mathbb{B} \mid \mathbb{A} \rightarrow \mathbb{B}\}$. Then $\text{CSP}(-, \mathbb{B})$ is solvable in polynomial time if and only if so is the problem $\text{SCSP}(\mathcal{C}_{\mathbb{B}}, -)$.*

PROOF. If the decision problem $\text{CSP}(-, \mathbb{B})$ is solvable in polynomial time we can construct an algorithm that given an instance $(\mathbb{A} \oplus \mathbb{B}, \mathbb{C})$ of $\text{SCSP}(\mathcal{C}_{\mathbb{B}}, -)$ computes a solution in polynomial time. The algorithm starts by computing an homomorphism φ from $\mathbb{A} \oplus \mathbb{B}$ to \mathbb{B} . Such a homomorphism exists by the definition of $\mathcal{C}_{\mathbb{B}}$ and can be computed in polynomial time because, by the aforementioned result of [11], if $\text{CSP}(-, \mathbb{B})$ is solvable in polynomial time then so is $\text{SCSP}(-, \mathbb{B})$. Then the algorithm decides by brute force whether or not there exists a homomorphism φ' from \mathbb{B} to \mathbb{C} (note that this can be done in polynomial time because \mathbb{B} is fixed). If such a homomorphism does not exist then we can certainly guarantee that there is no homomorphism from $\mathbb{A} \oplus \mathbb{B}$ to \mathbb{C} . Otherwise, the required homomorphism is obtained as the composition $\varphi' \circ \varphi$.

Conversely, assume that we have an algorithm Alg that finds a solution of any instance of $\text{CSP}(\mathcal{C}_{\mathbb{B}}, -)$ in polynomial time, say, $p(n)$. We construct from it an algorithm that solves $\text{CSP}(-, \mathbb{B})$. Given an instance (\mathbb{A}, \mathbb{B}) of $\text{CSP}(-, \mathbb{B})$ we call algorithm Alg with input $\mathbb{A} \oplus \mathbb{B}$ and \mathbb{B} . Additionally we count the number of steps performed by Alg in such a way that we stop if Alg has not finished in $p(n)$ steps. If Alg produces a correct answer then we have to be able to obtain from it a homomorphism from \mathbb{A} to \mathbb{B} . If Alg's answer is not correct or the clock reaches $p(n)$ steps we know that Alg failed. The only possible reason for that is that $\mathbb{A} \oplus \mathbb{B}$ does not belong to $\mathcal{C}_{\mathbb{B}}$, which implies that \mathbb{A} is not homomorphic to \mathbb{B} . \square

In what follows we transfer this result to enumeration problems. Let \mathcal{A} be a class of relational structures. The class \mathcal{A}' consists of all structures built as follows: Take $\mathbb{A} \in \mathcal{A}$ and add to it $|\mathbb{A}|$ independent vertices.

Lemma 2. *Let \mathcal{A} be a class of relational structures. Then $\text{SCSP}(\mathcal{A}, -)$ is solvable in polynomial time if and only if $\text{ECSP}(\mathcal{A}', -)$ is solvable w.p.d.*

PROOF. If $\text{ECSP}(\mathcal{A}, -)$ is enumerable w.p.d., then for any structure $\mathbb{A}' \in \mathcal{A}'$ it takes time polynomial in $|\mathbb{A}'|$ to find the first solution. Since \mathbb{A}' is only twice of the size of the corresponding structure \mathbb{A} , it takes only polynomial time to solve $\text{SCSP}(\mathcal{A}, -)$.

Conversely, given a structure $\mathbb{A}' = \mathbb{A} \cup I \in \mathcal{A}'$, where $\mathbb{A} \in \mathcal{A}$ and I is the set of independent elements, and any structure \mathbb{B} . The first homomorphism from \mathbb{A}' to \mathbb{B} can be found in polynomial time, since $\text{SCSP}(\mathcal{A}, -)$ is polynomial time solvable and the independent vertices can be mapped arbitrarily. Let the restriction of this homomorphism onto \mathbb{A} be φ . Then while enumerating all possible $|\mathbb{B}|^{|\mathbb{A}|}$ extensions of φ we buy enough time to enumerate all homomorphisms from \mathbb{A} to \mathbb{B} using brute force. \square

4. Sequence of bounded width endomorphisms

In this section we show that for every fixed k , all the homomorphisms from \mathbb{A} to \mathbb{B} can be enumerated with polynomial delay if a sequence of width k endomorphisms of \mathbb{A} is given in the input. Throughout this section, we consider a fixed sequence A_0, \dots, A_n and $\varphi_1, \dots, \varphi_n$ as in the definition of a sequence of width k endomorphisms (Definition 1). For brevity, we denote $\mathbb{A}[A_i]$ by \mathbb{A}_i .

We will enumerate the homomorphisms from \mathbb{A} to \mathbb{B} by first enumerating the homomorphisms from $\mathbb{A}_n, \mathbb{A}_{n-1}, \dots$ to \mathbb{B} and then transforming them to homomorphisms from \mathbb{A} to \mathbb{B} using the homomorphisms φ_i . We obtain the homomorphisms from \mathbb{A}_i by extending the homomorphisms from \mathbb{A}_{i+1} to the set $A_i \setminus A_{i+1}$; Lemma 3 below will be useful for this purpose. In order to avoid producing a homomorphism multiple times, we need a delicate classification (see Definitions 2 and 3 for the notions elementary homomorphisms and index of a homomorphism).

Lemma 3. *Let \mathbb{A}, \mathbb{B} be relational structures and $X_1 \subseteq X_2 \subseteq A$ subsets, and let g_0 be a homomorphism from $\mathbb{A}[X_1]$ to \mathbb{B} . For every fixed k , there is a polynomial-time algorithm $\text{HOMOMORPHISM-EXT}(\mathbb{A}, \mathbb{B}, X_1, X_2, g_0)$ that decides whether g_0 can be extended to a homomorphism from $\mathbb{A}[X_2]$ to \mathbb{B} , if the tree width of the induced substructure $\mathbb{A}[X_2 \setminus X_1]$ is at most k .*

PROOF. Let $Y = X_2 \setminus X_1$. We construct a structure \mathbb{Y} and an expansion \mathbb{B}^* of \mathbb{B} in such a way that Gaifman graph of \mathbb{Y} equals $G(\mathbb{A}[Y])$ and there is a homomorphism from \mathbb{Y} to \mathbb{B}^* if and only if there is one from $\mathbb{A}[X_2]$ to \mathbb{B} extending g_0 . Since $G(\mathbb{A}[Y])$ has tree width k , this can be checked in polynomial time.

For each $R \in \tau$, say, ℓ -ary, and each $\mathbf{a} = (a_1, \dots, a_\ell) \in R^{\mathbb{A}}$ such that $\{a_1, \dots, a_\ell\} \cap Y \neq \emptyset$, we introduce a new relational symbol $R_{\mathbf{a}}$ as follows. Let $(a_{i_1}, \dots, a_{i_m})$ be the list of all elements from $\{a_1, \dots, a_\ell\} \cap Y$ where $i_1 < \dots < i_m$ and for some $i_s \neq i_t$ it may happen that $a_{i_s} = a_{i_t}$. Then $R_{\mathbf{a}}$ is m -ary, it is interpreted on \mathbb{Y} as $R_{\mathbf{a}}^{\mathbb{Y}} = \{(a_{i_1}, \dots, a_{i_m})\}$, and it is interpreted on \mathbb{B}^* as

$$R_{\mathbf{a}}^{\mathbb{B}^*} = \{(b_{i_1}, \dots, b_{i_m}) \mid (b_1, \dots, b_\ell) \in R^{\mathbb{B}} \text{ and } b_j = g_0(a_j) \text{ for } a_j \in X_1\}.$$

In a sense, relations $R_{\mathbf{a}}$ describe all possible restrictions that the fixed values for the elements from X_1 impose on possible values for elements from Y .

It is straightforward that a homomorphism from \mathbb{Y} to \mathbb{B}^* exists if and only if there exists a homomorphism from $\mathbb{A}[X_2]$ to \mathbb{B} extending g_0 . Indeed, the restriction of any homomorphism $\mathbb{A}[X_2]$ to \mathbb{B} extending g_0 onto Y is a homomorphism from \mathbb{Y} to \mathbb{B}^* . Conversely, if φ is a homomorphism from \mathbb{Y} to \mathbb{B}^* then $g_0 \cup \varphi$ is a homomorphism of $\mathbb{A}[X_2]$ to \mathbb{B} . Finally, the Gaifman graph of \mathbb{Y} equals $G(\mathbb{A}[Y])$. \square

Definition 2. The *index* of a homomorphism φ from \mathbb{A} to \mathbb{B} is the largest t such that φ can be written as $\varphi = \psi \circ \varphi_t \circ \dots \circ \varphi_1$ for some homomorphism ψ from \mathbb{A}_t to \mathbb{B} . In particular, if φ cannot be written as $\varphi = \psi \circ \varphi_1$, then the index of φ is 0.

Observe that if the index of φ is at least t , then there is a unique ψ such that $\varphi = \psi \circ \varphi_t \circ \dots \circ \varphi_1$: This follows from the fact that $\varphi_t \circ \dots \circ \varphi_1$ is a surjective mapping from A to A_t , thus if ψ' and ψ'' differ on A_t , then $\psi' \circ \varphi_t \circ \dots \circ \varphi_1$ and $\psi'' \circ \varphi_t \circ \dots \circ \varphi_1$ differ on A .

Definition 3. A homomorphism ψ from \mathbb{A}_t to \mathbb{B} is *elementary*, if it cannot be written as $\psi = \psi' \circ \varphi_{t+1}$. A homomorphism is *reducible* if it is not elementary.

Lemma 4. If a homomorphism ψ from \mathbb{A}_t to \mathbb{B} is elementary, then $\varphi = \psi \circ \varphi_t \circ \dots \circ \varphi_1$ has index exactly t . Conversely, if homomorphism φ from \mathbb{A} to \mathbb{B} has index t and can be written as $\varphi = \psi \circ \varphi_t \circ \dots \circ \varphi_1$, then the homomorphism ψ from \mathbb{A}_t to \mathbb{B} is elementary.

PROOF. By definition, $\varphi = \psi \circ \varphi_t \circ \dots \circ \varphi_1$ has index at least t . If φ has index at least $t + 1$, then $\varphi = \psi' \circ \varphi_{t+1} \circ \varphi_t \circ \dots \circ \varphi_1$. By the uniqueness of the ψ , we have $\psi = \psi' \circ \varphi_{t+1}$, contradicting the fact that ψ is elementary. Thus the index of φ is exactly t .

For the second part, suppose that ψ is not elementary, i.e., $\psi = \psi' \circ \varphi_{t+1}$ for some homomorphism ψ' from \mathbb{A}_{t+1} to \mathbb{B} . Now $\varphi = \psi' \circ \varphi_{t+1} \circ \varphi_t \circ \dots \circ \varphi_1$, thus the index of φ is at least $t + 1$. \square

Lemma 4 suggests a way of enumerating all the homomorphisms from \mathbb{A} to \mathbb{B} without repetitions: for $t = 0, \dots, n$, we enumerate all the elementary homomorphisms from \mathbb{A}_t to \mathbb{B} , and for each such homomorphism ψ , we compute $\varphi = \psi \circ \varphi_t \circ \dots \circ \varphi_1$. To this end, we need the following characterization of elementary homomorphisms:

Lemma 5. A homomorphism ψ from \mathbb{A}_t to \mathbb{B} is reducible if and only if

- (1) $\psi(x) = \psi(y)$ for every $x, y \in A_t$ with $\varphi_{t+1}(x) = \varphi_{t+1}(y)$, i.e., for every $z \in A_{t+1}$, $\psi(x)$ has the same value b_z for every $x \in \varphi_{t+1}^{-1}(z)$, and
- (2) the mapping defined by $\psi'(z) := b_z$ is a homomorphism from \mathbb{A}_{t+1} to \mathbb{B} .

PROOF. Suppose first that both conditions hold. Then $\psi = \psi' \circ \varphi_{t+1}$ (where ψ' is as defined in the second condition). Since ψ' is a homomorphism from \mathbb{A}_{t+1} to \mathbb{B} , this means that ψ reducible.

Next we show that if ψ reducible, then both conditions hold. Suppose that $\psi = \psi'' \circ \varphi_{t+1}$, where ψ'' is a homomorphism from \mathbb{A}_{t+1} to \mathbb{B} . If there are two elements x, y such that $\varphi_{t+1}(x) = \varphi_{t+1}(y)$ and $\psi(x) \neq \psi(y)$, then we have a contradiction as $\psi(x) = \psi''(\varphi_{t+1}(x)) = \psi''(\varphi_{t+1}(y)) = \psi(y)$. Since φ_{t+1} is onto A_{t+1} , the mapping ψ'' is the same as the mapping ψ' defined in the second condition. Thus ψ' is a homomorphism from \mathbb{A}_{t+1} to \mathbb{B} . \square

Lemma 5 gives a way of testing in polynomial time whether a given homomorphism ψ is elementary: we have to test whether one of the two conditions are violated. We state this in a more general form: we can test in polynomial time whether a partial

mapping g_0 can be extended to an elementary homomorphism ψ , if the structure induced by the elements where g_0 is not defined has bounded tree width (see Lemma 7). We fix values every possible way in which the conditions of Lemma 5 can be violated and use HOMOMORPHISM-EXT to check whether there is an extension compatible with this choice. That is, we fix values every possible way that forces a violation of one of the two conditions in Lemma 5 and then we check whether g_0 can be extended in a way compatible with these fixed values. For example, to enforce the violation of (1) in Lemma 5, we need choose x and y with $\varphi_{t+1}(x) = \varphi_{t+1}(y)$ and fix different values for $\psi(x)$ and $\psi(y)$. To enforce the violation of condition (2), the obvious thing to do is to select a relation $R^{\mathbb{A}_{t+1}}$ of \mathbb{A}_{t+1} , a tuple $\mathbf{a} \in R^{\mathbb{A}_{t+1}}$, a tuple $\mathbf{b} \notin R^{\mathbb{B}}$, and fix values such that ψ' maps \mathbf{a} to \mathbf{b} . However, this would require going through all tuples \mathbf{b} *not appearing* in a relation $R^{\mathbb{B}}$. We follow a somewhat different approach to enumerate the possible violations more efficiently. We need the following definition:

Definition 4. Given a relation $R^{\mathbb{B}}$ of arity r , a *bad prefix* is a tuple $(b_1, \dots, b_s) \in B^s$ with $s \leq r$ such that

1. there is no tuple $(b_1, \dots, b_s, b_{s+1}, \dots, b_r) \in R^{\mathbb{B}}$ for any $b_{s+1}, \dots, b_r \in B$, and
2. there is a tuple $(b_1, \dots, b_{s-1}, c_s, c_{s+1}, \dots, c_r) \in R^{\mathbb{B}}$ for some $c_s, \dots, c_r \in B$.

If $(b_1, \dots, b_r) \notin R^{\mathbb{B}}$, then there is a unique $1 \leq s \leq r$ such that the tuple (b_1, \dots, b_s) is a bad prefix: there has to be an s such that (b_1, \dots, b_s) cannot be extended to a tuple of $R^{\mathbb{B}}$, but (b_1, \dots, b_{s-1}) can.

Lemma 6. *The relation $R^{\mathbb{B}}$ has at most $|R^{\mathbb{B}}| \cdot (|B| - 1) \cdot r$ bad prefixes, where r is the arity of the relation.*

PROOF. By definition, for every bad prefix (b_1, \dots, b_s) , there is a tuple

$$(b_1, \dots, b_{s-1}, c_s, c_{s+1}, \dots, c_r) \in R^{\mathbb{B}}.$$

Fix such a tuple for each bad prefix. Let us count how many bad prefixes are assigned to a tuple in $R^{\mathbb{B}}$. At most $|B| - 1$ bad prefixes of length s can be associated with a tuple: the bad prefix has to agree on the first $s - 1$ coordinates, and it has to be different on the s -th coordinate. Therefore, the total number of bad prefixes is at most $|R^{\mathbb{B}}| \cdot (|B| - 1) \cdot r$. \square

Lemma 7. *Let X be a subset of A_t and let g_0 be a mapping from X to B . For every fixed k , there is a polynomial-time algorithm ELEMENTARY-EXT(t, X, g_0) that decides whether g_0 can be extended to an elementary homomorphism from \mathbb{A}_t to B , if the tree width of the structure induced by $A_t - X$ is at most k .*

PROOF. We try to find a homomorphism that violates one of the conditions in Lemma 5. In order to do so, we try every possible way in which the conditions can be violated. First, we enumerate every possible way condition (1) can be violated. For this purpose, we enumerate every quadruple (x_1, x_2, b_1, b_2) with $x_1, x_2 \in A_t$, $\varphi_{t+1}(x_1) = \varphi_{t+1}(x_2)$, $b_1, b_2 \in B$, and $b_1 \neq b_2$. We try to find an extension of g_0 with $g_0(x_1) = b_1$ and $g_0(x_2) = b_2$; it is clear that if such an extension exists, then it is an elementary

homomorphism from \mathbb{A}_t to \mathbb{B} . If $x_1 \in X$ and $g_0(x_1) \neq b_1$, then such an extension does not exist (and similarly for x_2). Otherwise we can set $X' = X \cup \{x_1, x_2\}$ and extend g_0 by defining $g_0(x_1) = b_1$ and $g_0(x_2) = b_2$ (if it is not already defined so). Now we can apply Algorithm HOMOMORPHISM-EXT($\mathbb{A}, \mathbb{B}, X', A_t, g_0$) to check if g_0 can be extended from X' to A_t .

Next we try to find an extension that satisfies the first condition of Lemma 5 but violates the second. If ψ' is not a homomorphism, then there is a relation $R \in \tau$ and a tuple $\mathbf{a} = (z_1, \dots, z_r) \in R^{\mathbb{A}}$ with $z_1, \dots, z_r \in A_{t+1}$ such that $\psi'(\mathbf{a}) \notin R^{\mathbb{B}}$. We enumerate every $R \in \tau$, tuple $\mathbf{a} \in R^{\mathbb{A}} \cap A_{t+1}^r$, and every bad prefix (b_1, \dots, b_s) of $R^{\mathbb{B}}$. Let x_i be an arbitrary element of A_t with $\varphi_{t+1}(x_i) = z_i$. We extend g_0 by defining $g_0(x_i) = b_i$ for every $1 \leq i \leq s$. If $g_0(x_i)$ was already defined to have a different value, then we skip to the next bad prefix. Otherwise we get an extension of g_0 to $X' = X \cup \{x_1, \dots, x_s\}$. We show that if g_0 can be further extended from X' to a homomorphism ψ from \mathbb{A}_t to \mathbb{B} (which can be checked by calling HOMOMORPHISM-EXT($\mathbb{A}, \mathbb{B}, X', A_t, g_0$)), then this homomorphism ψ is an elementary homomorphism. Suppose that ψ does not violate (1) of Lemma 5 and let ψ' be as defined by the second condition. Since $\psi(x_i) = z_i$, we have that $\psi'(z_i) = \psi(x_i) = b_i$ for $1 \leq i \leq s$. Thus $(\psi'(z_1), \dots, \psi'(z_r)) \notin R^{\mathbb{B}}$ (since (b_1, \dots, b_s) is a bad prefix), which means that ψ' is not a homomorphism and the second condition is violated. Therefore, if g_0 has an elementary extension that satisfies the first condition and violates the second, then our algorithm finds an elementary extension when the appropriate relation R , tuple \mathbf{a} , and bad prefix (b_1, \dots, b_s) are considered. Thus we can conclude that algorithm ELEMENTARY-EXT(t, X', g_0) finds an elementary extension of g_0 if it exists. \square

We enumerate the elementary homomorphisms in a specific order defined by the following precedence relation.

Definition 5. Let φ be an elementary homomorphism from \mathbb{A}_i to \mathbb{B} and let ψ be an elementary homomorphism from \mathbb{A}_j to \mathbb{B} for some $j > i$. Homomorphism ψ is the *parent* of φ (φ is a *child* of ψ) if φ restricted to A_{i+1} can be written as $\psi \circ \varphi_j \circ \dots \circ \varphi_{i+2}$. *Ancestor* and *descendant* relations are defined as the reflexive transitive closure of the parent and child relations, respectively.

Note that an elementary homomorphism from \mathbb{A}_i to \mathbb{B} has exactly one parent for $i < n$ and a homomorphism from \mathbb{A}_n to \mathbb{B} has no parent. Fix an arbitrary ordering of the elements of A . For $0 \leq i \leq n$ and $0 \leq j \leq |A_i \setminus A_{i+1}|$, let $A_{i,j}$ be the union of A_{i+1} and the first j elements of $A_i \setminus A_{i+1}$. Note that $A_{i,0} = A_{i+1}$ and $A_{i,|A_i \setminus A_{i+1}|} = A_i$.

Lemma 8. Let ψ be a mapping from $A_{i,j}$ to B that can be extended to an elementary homomorphism from \mathbb{A}_i to \mathbb{B} . Assume that a sequence of width k endomorphisms is given for \mathbb{A} . For every fixed k , there is a polynomial-delay, polynomial-space algorithm ELEMENTARY-ENUM(i, j, ψ) that enumerates all the elementary homomorphisms from \mathbb{A}_i to \mathbb{B} that extends ψ and all the descendants of these homomorphisms.

PROOF. If $j < |A_i \setminus A_{i+1}|$, then we enumerate every element b of B , and extend ψ by defining $\psi'(a_{i,j+1}) = b$ and $\psi'(x) = \psi(x)$ for every $x \in A_{i,j}$. For every such ψ' , we

Input: Integers $0 \leq i \leq n$, $0 \leq j \leq |A_i \setminus A_{i+1}|$, mapping ψ from $A_{i,j}$ to \mathbb{B} .

Output: Every elementary homomorphism from \mathbb{A}_i to \mathbb{B} that extends ψ and all descendants of these homomorphisms.

```

Step 1  if  $j < |A_i \setminus A_{i+1}|$  then do
Step 1.1  for all  $b \in B$  do
Step 1.1.1  extend  $\psi$  to  $\psi'$  with  $\psi'(a_{i,j+1}) = b$ 
Step 1.1.2  if ELEMENTARY-EXT( $i, A_{i,j+1}, \psi'$ ) = true do
Step 1.1.2.1  ELEMENTARY-ENUM( $i, j + 1, \psi'$ )
Step 2  if  $j = |A_i \setminus A_{i+1}|$ 
Step 2.1  output  $\psi$ 
Step 2.2  for  $k := 1$  to  $i - 1$  then do
Step 2.2.1   $\psi_k := \psi \circ \varphi_i \circ \dots \circ \varphi_{k+2}$ 
Step 2.2.2  if ELEMENTARY-EXT( $k, A_{k+1}, \psi_k$ ) = true then do
Step 2.2.2.1  ELEMENTARY-ENUM( $k, 0, \psi_k$ )

```

Figure 1: Algorithm ELEMENTARY-ENUM(i, j, ψ)

use Algorithm ELEMENTARY-EXT($i, A_{i,j+1}, \psi'$) of Lemma 7 to check whether this extension ψ' can be further extended to an elementary homomorphism from \mathbb{A}_i to \mathbb{B} . If so, then we recursively call ELEMENTARY-ENUM($i, j + 1, \psi'$). Note that by the assumption that ψ has an extension to an elementary homomorphism from \mathbb{A}_i to \mathbb{B} , at least one choice of $b \in B$ results in a recursive call.

If $j = |A_i \setminus A_{i+1}|$ (which means that $A_{i,j} = A_i$), then ψ is an elementary homomorphism \mathbb{A}_i from \mathbb{B} , which we output. For every $1 \leq k \leq i - 1$, let $\psi_k = \psi \circ \varphi_i \circ \dots \circ \varphi_{k+2}$ be a mapping from A_{k+1} (i.e., $A_{k,0}$) to B . It is clear from the definition that if an elementary homomorphism φ of \mathbb{A}_k is a child of ψ , then φ extends ψ_k . For every $1 \leq k \leq i - 1$, we call ELEMENTARY-EXT($k, 0, \psi_k$) of Lemma 7 to check if ψ_k can be extended to an elementary homomorphism from \mathbb{A}_k to \mathbb{B} , and if so, then we make a recursive call ELEMENTARY-ENUM($k, 0, \psi_k$). It is clear that these recursive calls enumerate every child (and therefore every descendant) of ψ . Furthermore, as the different recursive calls enumerate different children (since k is different in each call), each descendant is enumerated exactly once.

Observe that the recursion depth is $O(|A|)$, the time spent at each node of the recursion tree is polynomial and we output an elementary homomorphism at every leaf node (a leaf node is possible only if $j = |A_i \setminus A_{i+1}|$). Thus the delay between two outputs is polynomial and the space requirement is also polynomial. \square

By calling ELEMENTARY-ENUM($n, 0, g_0$) (where g_0 is a trivial mapping from \emptyset to \mathbb{B}), we can enumerate all the elementary homomorphisms. By the observation in Lemma 4, this means that we can enumerate all the homomorphisms from \mathbb{A} to \mathbb{B} .

Theorem 1. *For every fixed k , there is a polynomial-delay, polynomial-space algorithm that, given structures \mathbb{A} , \mathbb{B} , and a sequence of width k endomorphisms of \mathbb{A} , enumerates all the homomorphisms from \mathbb{A} to \mathbb{B} .*

The following example shows that Theorem 1 does not provide a complete description of classes of structures solvable w.p.d.

Example 2. Let \mathcal{A} be the class of structures that are the disjoint union of a loop and a core. Obviously, $\text{SCSP}(\mathcal{A}, -)$ is polynomial time solvable. Recall that we denoted by \mathcal{A}' the class of all structures built by taking a $\mathbb{A} \in \mathcal{A}$ and adding to it $|\mathbb{A}|$ independent vertices. By Lemma 2, $\text{ECSP}(\mathcal{A}', -)$ is solvable with polynomial delay. However, it is not hard to see that \mathcal{A}' does not have a sequence of endomorphisms of bounded tree width.

Furthermore, as we will see in the next section it is hard, in general, to find a sequence of bounded width endomorphisms. Still, we can find a sequence of endomorphisms for a structure \mathbb{A} if we impose additional restrictions on the sequence. This is done in Section 6.

5. Hardness results

The first result of this section shows that finding a sequence of endomorphisms of bounded width can be difficult even in the simplest cases.

Theorem 2. *It is NP-complete to decide if a structure has a sequence of 1-width endomorphisms or a sequence of 1-width retractions to the core.*

PROOF. The proof is by reduction from 3SAT. Let ψ be a CNF formula with n variables and m clauses. We construct a relational structure \mathbb{A} (a colored graph) whose core has tree width 1. We show that \mathbb{A} has a sequence of endomorphisms to the core if and only if \mathbb{A} has a sequence of retractions if and only if ψ is satisfiable.

Construction. The core of \mathbb{A} has 6 nodes called $r, t, f, 1, 2, 3$ (see Figure 2). Vertex r has a self-loop and is connected to every other vertex of the core. Using distinct colors on the vertices of the core, we can ensure that this structure is indeed a core (in fact that the identity is its only endomorphism) and that the core is unique.




Figure 2: The structure \mathbb{A} constructed in the reduction.

Let us build a tree T the following way. There is a distinguished vertex named r' that will be called the root of the tree. This node is connected with n nodes v_i

($i = 1, \dots, n$) and m nodes c_j ($j = 1, \dots, m$). Each node v_i ($i = 1, \dots, n$) is connected to two descendants x_i and \bar{x}_i . Also we add for every node c_j ($j = 1, \dots, m$) three new nodes $c_{j,\ell}$ ($\ell = 1, 2, 3$) to which it is connected. We add colors to the nodes in T so that every node can only be mapped either to itself or to the core. This enforces that every endomorphism of \mathbb{A} is also a retraction. Also, by adding appropriately colors we can place some restrictions on to which element of the core a given element of T can be mapped. In particular, nodes v_i , $i = 1, \dots, n$ can only be mapped to t or f , nodes c_j , $j = 1, \dots, m$ can only be mapped to 1, 2 or 3, and the rest of nodes of T can only be mapped to r .

We add some additional edges connecting the leaves of T , thus T will no longer be a tree. These edges encode the structure of the formula ψ : if the ℓ -th literal of the j -th clause is the literal x_i (resp., \bar{x}_i), then connect $c_{j,\ell}$ with \bar{x}_i (resp., x_i).

To complete the description of the structure, we define the connections between the core and T . Vertex f is connected with each x_i ($i = 1, \dots, n$) whereas vertex t is connected with each \bar{x}_i ($i = 1, \dots, n$). Each vertex $c_{j,\ell}$ ($\ell = 1, 2, 3$) is connected with exactly two of vertices 1, 2, 3 of the core: in particular it is not connected to vertex ℓ but connected to the other two. Finally, r is connected to each x_i, \bar{x}_i ($i = 1, \dots, n$), $c_{j,\ell}$ ($j = 1, \dots, m, \ell = 1, 2, 3$).

Endomorphisms \Rightarrow assignment of ψ . Assume that \mathbb{A} has a sequence of 1-width endomorphisms to the core. Let φ be the first endomorphism, which, as we observed before, must be a retraction.

Assume that φ maps some vertex v of T to the core. Notice that if a vertex v of T is mapped to the core, then the parent of v is also mapped to the core: this follows from the fact that vertices $v_1, \dots, v_n, c_1, \dots, c_m, r'$ have no connections to the core. Therefore, we can assume that the root vertex r' of T is mapped to the core, in particular to r . As every descendant of r' is not connected to r , it follows that it must be mapped to the core. Hence every node v_i ($i = 1, \dots, n$) is mapped either to t or f and every c_j ($j = 1, \dots, m$) is mapped to either 1, 2 or 3.

Define an assignment of ψ by setting variable x_i to true if and only if v_i is mapped to t . We claim that this is a satisfying assignment. For every $j = 1, \dots, m$, let ℓ be the node in the core to which c_j is mapped. We claim that the ℓ -th literal of the j -th clause is true in the assignment and hence the clause is satisfied. Assume first that the ℓ -th literal is the positive literal x_i . If x_i was assigned the value false, then this means v_i is mapped to f . As f is not connected to \bar{x}_i , necessarily \bar{x}_i is mapped to the core. Similarly, if c_j is mapped to ℓ it follows that $c_{j,\ell}$ is mapped to r . By construction \bar{x}_i and $c_{j,\ell}$ are connected, which creates the following cycle in the vertices mapped to the core: $r', v_i, \bar{x}_i, c_{j,\ell}, c_j$, contradicting the assumption that the vertices mapped to the core induce a graph with tree width 1. In a similar way, if the ℓ -th literal is \bar{x}_i , then vertex x_i is mapped to the core, again creating a cycle.

Assignment of $\psi \Rightarrow$ retractions. Assume that ψ has a satisfying assignment. We construct a retraction φ_1 as follows. If x_i is true (resp., false) in the assignment, then we map vertex x_i (resp., \bar{x}_i) to r and we map its ancestor v_i to t (resp., f). For every j , there is an $1 \leq \ell \leq 3$ such that the ℓ -th literal of the j -th clause is true. For every such j and ℓ , vertex $c_{j,\ell}$ is mapped to r and vertex c_j is mapped to ℓ . Furthermore vertex

r' is mapped to r . From the fact that the assignment is satisfying, it follows that the leaves of T that are mapped to the core are independent. This means that the vertices in $\mathbb{A} - \varphi_1(\mathbb{A})$ induce a graph with tree width 1.

After applying retraction φ_1 , the vertices outside the core are of the form x_i, \bar{x}_i , or $c_{j,\ell}$. These vertices induce a set of stars and independent vertices (since the degree of every vertex $c_{j,\ell}$ is at most 1), thus they induce a graph with tree width at most 1. Therefore, we can map these vertices to r by a single retraction. \square

The second result shows that $\text{ECSP}(\mathcal{A}, -)$ can be hard even if every structure in \mathcal{A} has a sequence of width-2 endomorphisms. Note that this result is incomparable with Theorem 2, since an enumeration algorithm (in theory) does not necessarily have to compute a sequence of endomorphisms. We need the following lemma:

Lemma 9. *If G is a planar graph, then it is possible to find a partition (V_1, V_2) of its vertices in polynomial time such that $G[V_1]$ and $G[V_2]$ have tree width at most 2.*

PROOF. A planar embedding of G can be found in polynomial time [19]; let us fix a planar embedding of G . Define the *level* of a vertex as follows: vertices of the outer face have level 1, and a vertex is on level ℓ for some $\ell > 1$ if it is on the outer face after deleting every vertex of level less than ℓ . Observe that the level numbers of adjacent vertices differ by at most 1. Let V_1 (resp., V_2) be the vertices with odd (resp., even) level number. A connected component of $G[V_1]$ contains vertices with the same level number, which means that this component is outerplanar: the embedding of G gives an embedding of $G[V_1]$ where every vertex is on the infinite face of the embedding. Thus $G[V_1]$ (and similarly, $G[V_2]$) is outerplanar, and it is well known that every outerplanar graph has tree width at most 2 (cf. [5]). \square

Proposition 1. *There is a class \mathcal{A} of relational structures such that every structure from \mathcal{A} has a sequence of width 2 endomorphisms to the core, and such that the problem $\text{ECSP}(\mathcal{A}, -)$ is not solvable w.p.d., unless $P = NP$.*

PROOF. Let \mathcal{A} be a class of graphs built in the following way. Take a 3-colorable planar graph G and its partition (V_1, V_2) according to Lemma 9. Using colorings we can ensure that G is a core. Then we take a disjoint union of this graph with a triangle T having all the colors and a copy G_1 of $G[V_1]$. Let \mathbb{A} denote the resulting structure.

CLAIM 1. \mathbb{A} has a sequence of width-2 endomorphisms.

Let ψ be a 3-coloring of G that is a homomorphism into the triangle, and ψ' the bijective mapping from G_1 to $G[V_1]$. Then φ_1 is defined to act as ψ on G , as ψ' on G_1 and identically on T . Endomorphism φ_2 is just the 3-coloring of $G \cup G_1$ induced by ψ . The images of φ_1 and φ_2 are $T \cup G[V_1]$ and T , respectively, so all the conditions on a sequence of width-2 homomorphisms are easily checkable.

CLAIM 2. The PLANAR GRAPH 3-COLORING PROBLEM is polynomial-time Turing reducible to $\text{ECSP}(\mathcal{A}, -)$.

Given a planar graph G we find its partition (V_1, V_2) and create a structure \mathbb{A} , as described above. Then we apply an algorithm that enumerates solutions to $\text{ECSP}(\mathcal{A}, -)$. We may assume that such an algorithm stops with some time bound regardless whether G is 3-colorable or not. If the algorithm succeeds we can now produce a 3-coloring of G . \square

6. Finite extensions

We can find a sequence of endomorphisms for a structure \mathbb{A} if we impose two more restrictions on such a sequence.

A retraction φ of a structure \mathbb{A} is called a *k-retraction* if at most k nodes change their value according to φ . A structure is a *k-core* if the only k -retraction is the identity. A k -core of a structure is any k -core obtained by a sequence of k -retractions.

Let \mathbb{A} be a structure and let $B, B' \subseteq A$. We say that B and B' are \mathbb{A} -*identical* if there exists endomorphisms φ, φ' of \mathbb{A} such that $\varphi(B) = B'$ and $\varphi'(B') = B$. Notice that the definition implies that $\mathbb{A}[B]$ and $\mathbb{A}[B']$ are isomorphic.

Lemma 10. *Let \mathbb{A} be a structure, let φ be a k -retraction, and let ψ be a retraction (not necessarily a k -retraction) such that its image $\psi(\mathbb{A}) = \mathbb{B}$ is a k -core. Then B and $\varphi(B)$ are \mathbb{A} -identical.*

PROOF. Let \mathbb{B}' be the substructure of \mathbb{B} containing all $b \in \mathbb{B}$ such that $\varphi(b) = b$. Observe that there are at most k elements in $\mathbb{B} - \mathbb{B}'$. Now consider the mapping $\psi \circ \varphi$. This mapping acts as the identity on \mathbb{B}' . Furthermore, it sends every element of $\mathbb{B} - \mathbb{B}'$ to some element of \mathbb{B} . Consequently the restriction χ of $\psi \circ \varphi$ to \mathbb{B} is an endomorphism of \mathbb{B} which acts as the identity on \mathbb{B}' . Indeed, χ has to be an automorphism. To see it, notice that otherwise we could find a power of χ , $\chi^n = \chi \circ \dots \circ \chi$, that would be a proper retraction and since χ^n must act as the identity on any element of \mathbb{B}' it would contradict the fact that \mathbb{B} is a k -core. Consequently, φ and ψ certify that B and $\varphi(B)$ are \mathbb{A} -identical. \square

Lemma 11. *All k -cores of a structure \mathbb{A} are isomorphic.*

PROOF. Let \mathbb{B} and \mathbb{C} be two k -cores obtained following different sequences of k -retractions. Let $\varphi_1, \dots, \varphi_n$ be the sequence of k -retractions that produces \mathbb{C} . and for $i = 1, \dots, n$ define φ'_i to be $\varphi_i \circ \dots \circ \varphi_1$. We prove by induction that

$$(*) \quad \varphi'_i(B) \text{ and } B \text{ are } \mathbb{A}\text{-identical.}$$

The case $i = 1$ can be solved just by assuming that φ_1 is the identity mapping. For the inductive step we need to prove that $\varphi'_{i+1}(B)$ and $\varphi'_i(B)$ are \mathbb{A} -identical. In order to do this we apply Lemma 10. We just need to find a retraction of \mathbb{A} whose image is $\varphi'_i(B)$. By inductive hypothesis there exists endomorphisms χ and χ' of \mathbb{A} such that $\chi(B) = \varphi'_i(B)$ and $\chi'(\varphi'_i(B)) = B$. Consider the mapping $\chi \circ \psi \circ \chi'$ where ψ is the retraction with image B given by the hypothesis of the Lemma. We have $\chi \circ \psi \circ \chi'(A) = \chi \circ \psi \circ \chi'(\varphi'_i(B)) = \varphi'_i(B)$. Consequently, some power of this mapping gives the desired retraction. This finishes the proof of (*).

We have just seen that \mathbb{C} has a substructure $\varphi'_n(\mathbb{B})$ which is isomorphic to \mathbb{B} . By a symmetric argument we conclude that \mathbb{B} contains as a substructure an isomorphic copy of \mathbb{C} . Hence \mathbb{B} and \mathbb{C} are isomorphic. \square

The following result follows from Lemma 11 and Theorem 1.

Theorem 3. *Let $k > 0$ be a positive integer and let \mathcal{C} be a class of structures such that the k -core of every structure in \mathcal{C} has tree width at most k . Then, the enumeration problem $\text{ECSP}(\mathcal{C}, -)$ is solvable w.p.d.*

PROOF. From Lemma 11 it follows that given an instance (\mathbb{A}, \mathbb{B}) of $\text{ECSP}(\mathcal{C}, -)$ it is possible to compute by greedy search the k -core of \mathbb{A} along with the sequence of k -retractions leading to it. Since every sequence of k -retractions is also a sequence of width k endomorphisms it follows from Theorem 1 that one can enumerate in polynomial time all homomorphisms from \mathbb{A} to \mathbb{B} .

□

Corollary 1. *If \mathcal{C} is a class of structures of bounded tree width then $\text{ECSP}(\mathcal{C}, -)$ is solvable w.p.d.*

7. Conjunctive queries

When making a query to a database the user usually needs to obtain values of only those variables (attributes) (s)he is interested in. In terms of homomorphisms this can be translated as follows: For relational structures \mathbb{A}, \mathbb{B} , and a subset $Y \subseteq A$, we aim to list those mappings from Y to B which can be extended to a full homomorphism from \mathbb{A} to \mathbb{B} . In other words, we would like to enumerate all the mappings from Y to B that arise as the restriction of some homomorphism from \mathbb{A} to \mathbb{B} . Clearly, this problem significantly differs from the regular enumeration problem. A mapping from Y to B can be extendible to a homomorphism in many ways, possibly superpolynomially many, and an enumeration algorithm would list all of them. In the worst case scenario it would list them before turning to the next partial mapping. If this happens it may destroy polynomiality of the delay between outputting consecutive solutions.

In this section we treat the CONJUNCTIVE QUERY EVALUATION PROBLEM as follows.

$\text{CQE}(\mathcal{A}, \mathcal{B})$
Instance: $\mathbb{A} \in \mathcal{A}, \mathbb{B} \in \mathcal{B}, Y \subseteq A$
Problem: Output all partial mappings from Y to B extendible to a homomorphism from \mathbb{A} to \mathbb{B} .

It follows from [16] that if a class \mathcal{A} of bounded arity does not have bounded tree width modulo homomorphic equivalence then $\text{CQE}(\mathcal{A}, -)$ is not solvable w.p.d., unless $\text{FPT} = \text{W}[1]$. We present two new results about computing the solutions w.p.d. The first one shows that the problem $\text{CQE}(\mathcal{A}, -)$ is solvable w.p.d. if \mathcal{A} is a class of structures of bounded tree width. The second one claims that, modulo some complexity assumptions, in contrast to enumeration problems this cannot be generalized to structures with k -cores of bounded tree width for $k \geq 2$.

Theorem 4. *If \mathcal{A} is a class of structures of bounded width then $\text{CQE}(\mathcal{A}, -)$ is solvable w.p.d.*

PROOF. Let $(\mathbb{A}, \mathbb{B}, Y)$ be an instance of $\text{CQE}(\mathcal{A}, -)$. Fix arbitrary orderings on Y and B , which induce a natural lexicographic ordering on the partial mappings from Y to B . More precisely, let φ and ψ be any partial mappings from Y to B which we consider as mappings from Y to $B \cup \{\perp\}$ by setting every undefined element to \perp . Then we say that φ precedes lexicographically ψ if there exists some $y \in Y$, such that $\varphi(y) < \psi(y)$ and $\varphi(y') = \psi(y')$ for every $y' \in Y$ with $y' < y$ where $\perp < b$ for every $b \in B$.

It can be easily derived, from Lemma 3, a polynomial-time algorithm that computes, given a partial mapping φ from Y to B , the next partial partial mapping ψ in the lexicographical order that extends to an homomorphism from \mathbb{A} to \mathbb{B} or reports that such a partial mapping does not exist. To achieve this it is only necessary to compute the largest $y \in Y$ and smallest $b \in B$ with $b > \varphi(y)$ such that $\varphi_{y,b}$ can be extended to an homomorphism from \mathbb{A} to \mathbb{B} where $\varphi_{y,b}$ is the partial mapping defined as

$$\varphi_{y,b}(y') = \begin{cases} \varphi(y') & \text{if } y' < y \\ b & \text{if } y' = y \\ \perp & \text{if } y' > y \end{cases}$$

This can be achieved in polynomial time by Lemma 3 since the number of choices for (y, b) is polynomial.

Clearly, if such y and b exist then we can set ψ to $\varphi_{y,b}$ whereas otherwise we can conclude that no partial mapping larger than φ in the lexicographical order can be extended. Using this procedure one can derive an algorithm CQE-BOUNDED-WIDTH (in Figure 3) that outputs all solutions w.p.d. In a nutshell, algorithm CQE-BOUNDED-WIDTH computes in lexicographical order all partial mappings from Y to B that extend to an homomorphism from \mathbb{A} to \mathbb{B} and outputs only those that are defined over the whole Y . \square

Figure 3: Algorithm CQE-BOUNDED-WIDTH

Input: Relational structures \mathbb{A}, \mathbb{B} , and $Y = \{y_1, \dots, y_\ell\} \subseteq A$

Output: A list of mappings $\varphi: Y \rightarrow B$ extendible to a homomorphism from \mathbb{A} to \mathbb{B}

```

Step 1  set  $m = 0, \varphi = \emptyset, S_i = B, i \in [\ell], \text{complete} := \text{false}$ 
Step 2  while not complete do
Step 2.1 if  $m < \ell$  then do
Step 2.1.1 search  $S_{m+1}$  until a  $b \in S_{m+1}$  is found such that there exists a homomorphism extending
            $\varphi \cup \{y_{m+1} \rightarrow b\}$  and remove all members of  $S_{m+1}$  preceding  $b$  inclusive
Step 2.1.2 if such a  $b$  exists then set  $\varphi := \varphi \cup \{y_{m+1} \rightarrow b\}, m := m + 1$ 
Step 2.1.3 else
Step 2.1.3.1 if  $m \neq 0$  then set  $\varphi = \varphi|_{\{y_1, \dots, y_{m-1}\}}$  and  $S_{m+1} := B, m := m - 1$ 
Step 2.1.3.2 else set  $\text{complete} := \text{true}$ 
Step 2.2 else then do
Step 2.2.1 output  $\varphi$ 
Step 2.2.2 set  $\varphi := \varphi|_{\{y_1, \dots, y_{m-1}\}}, m := \ell - 1$ 
           endwhile

```

Theorem 4 does not generalize to classes of structures whose k -cores have bounded width.

Example 3. Recall that the MULTICOLORED CLIQUE problem (cf. [14]) is formulated as follows: Given a number k and a vertex k -colored graph, decide if the graph contains a k -clique all vertices of which are colored different colors. This problem is $W[1]$ -complete, i.e., has no time $f(k)n^c$ algorithm for any function f and constant c , unless $\text{FPT} = W[1]$. We reduce this problem to $\text{CQE}(\mathcal{A}, -)$ where \mathcal{A} is the class of structures whose 2-cores are 2-element described below.

Let us consider relational structures with two binary and two unary relations. This structure can be thought of as a graph whose vertices and edges have one of the two colors, say, red and blue, accordingly to which of the two binary/unary relations they belong to. Let \mathbb{A}_k be the relational structure with universe $\{a_1, \dots, a_k, y_1, \dots, y_k\}$, where a_1, \dots, a_k are red while y_1, \dots, y_k are blue. Then $\{a_1, \dots, a_k\}$ induces a red clique, that is every a_i, a_j (i, j are not necessarily different) are connected with a red edge, and each y_i is connected to a_i with a blue edge. It is not hard to see that every pair of a red and blue vertices induces a 2-core of this structure. Set $\mathcal{A} = \{\mathbb{A}_k \mid k \in \mathbb{N}\}$.

The reduction of the MULTICOLORED CLIQUE problem to $\text{CQE}(\mathcal{A}, -)$ goes as follows. Let $G = (V, E)$ be k -colored graph whose coloring induces a partition of V into classes B_1, \dots, B_k . Then we define structures \mathbb{A}, \mathbb{B} and a set $Y \subseteq A$. We set $\mathbb{A} = \mathbb{A}_k$, $Y = \{y_1, \dots, y_k\}$. Then let $B = V \cup \{b_1, \dots, b_k\}$, the elements of V are colored red and the induced substructure $\mathbb{B}[V]$ is the graph G (without coloring) whose edges are colored also red and in which we add a red loop to every node. Finally, b_1, \dots, b_k are made blue and each b_i is connected with a blue edge with every vertex from B_i .

It is not hard to see that any homomorphism maps $\{a_1, \dots, a_k\}$ to V and Y to $\{b_1, \dots, b_k\}$, and that the number of homomorphisms that do not agree on Y does not exceed k^k . Moreover, G contains a k -colored clique if and only if there is a homomorphism from \mathbb{A} to \mathbb{B} that maps Y onto $\{b_1, \dots, b_k\}$. If there existed an algorithm solving $\text{CQE}(\mathcal{A}, -)$ w.p.d., say, time needed to compute the first and every consequent solution is bounded by a polynomial $p(n)$, then time needed to list all solutions is at most $k^k p(n)$. This means that MULTICOLORED CLIQUE is FPT, a contradiction.

References

- [1] Albert Atserias, Martin Grohe, and Dániel Marx. Size bounds and query plans for relational joins. In *49th Annual IEEE Symposium on Foundations of Computer Science (FOCS'08)*, pages 739–748, 2008.
- [2] Per Austrin. Towards sharp inapproximability for any 2-CSP. *SIAM J. Comput.*, pages 2430–2466, 2010.
- [3] Libor Barto, Marcin Kozik, and Todd Niven. Graphs, polymorphisms and the complexity of homomorphism problems. *SIAM J. Comput.*, 38(5):1782–1802, 2009.
- [4] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *J. ACM*, 57(2), 2010.

- [5] H. L. Bodlaender. A tourist guide through treewidth. *Acta Cybernet.*, 11(1-2):1–21, 1993.
- [6] A. Bulatov and V. Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. *Information and Computation*, 205(5):651–678, 2007.
- [7] A. Bulatov and M. Grohe. The complexity of partition functions. *Theoretical Computer Science*, 348:148–186, 2005.
- [8] A.A. Bulatov. Tractable conservative constraint satisfaction problems. In *Proceedings of the 18th Annual IEEE Symposium on Logic in Computer Science*, pages 321–330, Ottawa, Canada, June 2003. IEEE Computer Society.
- [9] Andrei A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM*, 53(1):66–120, 2006.
- [10] Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. In *ICALP*, pages 646–661, 2008.
- [11] Andrei A. Bulatov, Peter Jeavons, and Andrei A. Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005.
- [12] M.E. Dyer, L.A. Goldberg, and M. Paterson. On counting homomorphisms to directed acyclic graphs. *Journal of the ACM*, 54(6), 2007.
- [13] T. Feder and M.Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM Journal of Computing*, 28:57–104, 1998.
- [14] Michael R. Fellows, Danny Hermelin, Frances A. Rosamond, and Stéphane Vialette. On the parameterized complexity of multiple-interval graph problems. *Theor. Comput. Sci.*, 410(1):53–61, 2009.
- [15] Gianluigi Greco and Francesco Scarcello. Structural tractability of enumerating CSP solutions. In *CP 1010*, pages 236–251, 2010.
- [16] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM*, 54(1), 2007.
- [17] P. Hell and J. Nešetřil. The core of graph. *Discrete Mathematics*, 109(1-3):117–126, 1992.
- [18] P. Hell and J. Nešetřil. On the complexity of H -coloring. *Journal of Combinatorial Theory, Ser.B*, 48:92–110, 1990.
- [19] John Hopcroft and Robert Tarjan. Efficient planarity testing. *J. Assoc. Comput. Mach.*, 21:549–568, 1974.

- [20] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. On generating all maximal independent sets. *Inf. Process. Lett.*, 27(3):119–123, 1988.
- [21] Peter Jonsson, Mikael Klasson, and Andrei A. Krokhn. The approximability of three-valued MAX CSP. *SIAM J. Comput.*, 35(6):1329–1349, 2006.
- [22] Peter Jonsson and Andrei A. Krokhn. Maximum H-colourable subdigraphs and constraint optimization with arbitrary weights. *J. Comput. Syst. Sci.*, 73(5):691–702, 2007.
- [23] Dániel Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. In *STOC*, pages 735–744, 2010.
- [24] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *STOC 2008*, pages 245–254, 2008.
- [25] Francesco Scarcello, Georg Gottlob, and Gianluigi Greco. Uniform constraint satisfaction problems and database theory. In *Complexity of Constraints 2008*, 2008.
- [26] Henning Schnoor and Ilka Schnoor. Enumerating all solutions for constraint satisfaction problems. In *STACS 2007*, pages 694–705, 2007.