The complexity of general-valued CSPs seen from the other side

Clément Carbonnel, Miguel Romero, Stanislav Živný
University of Oxford, UK
$\{$clement.carbonnel,miguel.romero,standa.zivny$\}$@cs.ox.ac.uk

***Abstract*—The constraint satisfaction problem (CSP) is concerned with homomorphisms between two structures. For CSPs with restricted left-hand side structures, the results of Dalmau, Kolaitis, and Vardi [CP'02], Grohe [FOCS'03/JACM'07], and Atserias, Bulatov, and Dalmau [ICALP'07] establish the precise borderline of polynomial-time solvability (subject to complexity-theoretic assumptions) and of solvability by bounded-consistency algorithms (unconditionally) as bounded treewidth modulo homomorphic equivalence.**

The general-valued constraint satisfaction problem (VCSP) is a generalisation of the CSP concerned with homomorphisms between two *valued* structures. For VCSPs with restricted left-hand side valued structures, we establish the precise borderline of polynomial-time solvability (subject to complexity-theoretic assumptions) and of solvability by the k-th level of the Sherali-Adams LP hierarchy (unconditionally). We also obtain results on related problems concerned with finding a solution and recognising the tractable cases; the latter has an application in database theory.

***Keywords*-valued constraint satisfaction; homomorphism problems; fractional homomorphism; treewidth; Sherali-Adams LP relaxation**

I. INTRODUCTION

Constraint Satisfaction Problems: The homomorphism problem for relational structures is a fundamental computer science problem: Given two relational structures $\mathbf{A}$ and $\mathbf{B}$ over the same signature, the goal is to determine the existence of a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ (see, e.g., the book by Hell and Nešetřil on this topic [31]). The homomorphism problem is known to be equivalent to the evaluation problem and the containment problem for conjunctive database queries [12], [33], and also to the constraint satisfaction problem (CSP) [21], which originated in artificial intelligence [39] and provides a common framework for expressing a wide range of both theoretical and real-life combinatorial problems.

For a class $\mathcal{C}$ of relational structures, we denote by CSP($\mathcal{C}$, $-$) the restriction of the homomorphism problem in which the input structure $\mathbf{A}$ belongs to $\mathcal{C}$ and the input structure $\mathbf{B}$ is arbitrary (these types of restrictions are known as *structural* restrictions). Similarly, by CSP($-$, $\mathcal{C}$) we denote the restriction of the homomorphism problem in which the input structure $\mathbf{A}$ is arbitrary and the input structure $\mathbf{B}$ belongs to $\mathcal{C}$.

Feder and Vardi initiated the study of CSP($-$, $\{\mathbf{B}\}$), also known as non-uniform CSPs, and famously conjectured that, for every fixed finite structure $\mathbf{B}$, either CSP($-$, $\{\mathbf{B}\}$) is in PTIME or CSP($-$, $\{\mathbf{B}\}$) is NP-complete. For example, if $\mathbf{B}$ is a clique on k vertices then CSP($-$, $\{\mathbf{B}\}$) is the well-known k-colouring problem, which is known to be in PTIME for $k \leq 2$ and NP-complete for $k \geq 3$. Most of the progress on the Feder-Vardi conjecture (e.g., [6], [3], [32], [10], [2]) is based on the algebraic approach [9], culminating in two recent (affirmative) solutions to the Feder-Vardi conjecture obtained independently by Bulatov [7] and Zhuk [47].

Note that CSP($\mathcal{C}$, $-$) is only interesting if $\mathcal{C}$ is an infinite class of structures as otherwise CSP($\mathcal{C}$, $-$) is always in PTIME. (This is, however, not the case for CSP($-$, $\mathcal{C}$) as we have seen in the example of 3-colouring.) Freuder observed that CSP($\mathcal{C}$, $-$) is in PTIME if $\mathcal{C}$ consists of trees [23] or, more generally, if it has bounded treewidth [24]. Later, Dalmau, Kolaitis, and Vardi showed that CSP($\mathcal{C}$, $-$) is solved by k-consistency, a fundamental local propagation algorithm [15], if $\mathcal{C}$ is of bounded treewidth *modulo homomorphic equivalence*, i.e., if the treewidth of the cores of the structures from $\mathcal{C}$ is at most k, for some fixed $k \geq 1$ [14]. Atserias, Bulatov, and Dalmau showed that this is precisely the class of structures solved by k-consistency [1]. In [29], Grohe proved that the tractability result of Dalmau et al. [14] is optimal for classes $\mathcal{C}$ of bounded arity: Under the assumption that FPT $\neq$ W[1], CSP($\mathcal{C}$, $-$) is tractable if and only if $\mathcal{C}$ has bounded treewidth modulo homomorphic equivalence.

General-valued Constraint Satisfaction Problems: General-valued Constraint Satisfaction Problems (VCSPs) are generalisations of CSPs which allow for not only decision problems but also for optimisation problems (and the mix of the two) to be considered in one framework [13]. In the case of VCSPs we deal with *valued* structures. Regarding tractable restrictions, the situation of the non-uniform case is by now well-understood. Indeed, assuming the (now proved) Feder-Vardi conjecture, it holds that for any fixed valued structure $\mathbb{B}$, either VCSP($-$, $\{\mathbb{B}\}$) is in PTIME or VCSP($-$, $\{\mathbb{B}\}$) is NP-complete [37], [35].

For structural restrictions, it is a folklore result that VCSP($\mathcal{C}$, $-$) is tractable if $\mathcal{C}$ is of bounded treewidth; see, e.g. [4]. So is the fact that the $(k+1)$-st level of the Sherali-Adams LP hierarchy [43] solves VCSP($\mathcal{C}$, $-$) to optimality if the treewidth of all structures in $\mathcal{C}$ is at most k. (We are not aware of any reference for this fact. For certain special problems, it is discussed in [5]. For the extension complexity of such problems, see [34].) However, unlike the CSP case, the precise borderline of polynomial-time solvability and the power of fundamental algorithms (as the Sherali-Adams LP hierarchy) for VCSP($\mathcal{C}$, $-$) is still unknown. Understanding these complexity and algorithmic frontiers for VCSP($\mathcal{C}$, $-$) is the main goal of this paper.

Contribution: We study the problem VCSP($\mathcal{C}$, $-$) for classes $\mathcal{C}$ of valued structures. As our first result, we give (in Theorem IV.1) a complete complexity classification of VCSP($\mathcal{C}$, $-$) and identify the precise borderline of tractability, for classes $\mathcal{C}$ of bounded arity. A key ingredient in our result is a novel notion of *valued equivalence* and a characterisation of this notion in terms of *valued cores*. More precisely, we show that VCSP($\mathcal{C}$, $-$) is tractable if and only if $\mathcal{C}$ has bounded treewidth modulo valued equivalence. This latter notion *strictly* generalises bounded treewidth and it is *strictly* weaker than bounded treewidth modulo homomorphic equivalence. Our proof builds on the characterisation by Dalmau et al. [14] and Grohe [29] for CSPs (see Section IV).

We show that the newly identified tractable classes are solvable by the Sherali-Adams LP hierarchy. Our second result (Theorem V.1) gives a precise characterisation of the power of Sherali-Adams for VCSP($\mathcal{C}$, $-$). In particular, we show that the $(k+1)$-st level of the Sherali-Adams LP hierarchy solves VCSP($\mathcal{C}$, $-$) to optimality if and only if the valued cores of the structures from $\mathcal{C}$ have treewidth *modulo scopes* at most k and the *overlaps* of scopes are of size at most $k+1$. The proof builds on the work of Atserias et al. [1] and Thapper and Živný [46], as well as on an adaptation of the classical result connecting treewidth and brambles by Seymour and Thomas [42].

Our main results are for the VCSP in which we ask for the cost of an optimal solution. It is also possible to define the VCSP as a *search* problem, in which one is additionally required to return a solution with the optimal cost. A complete characterisation of tractable search cases in terms of structural properties of (a class of structures) $\mathcal{C}$ is open even for CSPs and there is some evidence that the tractability frontier cannot be captured in simple terms [8, Lemma 1]. Building on our main results as well as on techniques from [45], we give in Section VI a characterisation of the tractable cases for search VCSP($\mathcal{C}$, $-$) in terms of tractable core computation (Theorem VI.1). Finally, we provide in Section VII tight complexity bounds for several problems related to our classification results, e.g., deciding whether the treewidth is at most k modulo valued equivalence, deciding solvability by the k-th level of the Sherali-Adams LP hierarchy, and deciding valued equivalence for valued structures. These results have interesting consequences to database theory, specifically, to the evaluation and optimisation of conjunctive queries over *annotated* databases. In particular, we show that the containment problem of conjunctive queries over the tropical semiring is in NP, thus improving on the work of [36], which put it in Π_2^p.

Related work: In his PhD thesis [19], Färnqvist studied the complexity of VCSP($\mathcal{C}$, $-$) and also some fragments of VCSPs (see also [20], [18]). He considered a very specific framework that only allows for particular types of classes $\mathcal{C}$'s to be classified. For these classes, he showed that only bounded treewidth gives rise to tractability (assuming bounded arity) and asked about more general classes. In particular, decision CSPs do *not* fit in his framework and Grohe's classification [29] is not implied by Färnqvist's work. In contrast, our characterisation (of *all* classes $\mathcal{C}$'s of valued structures) gives rise to new tractable cases going beyond those identified by Färnqvist. Moreover, we can derive both Grohe's classification and Färnqvist's classification directly from our results, as explained in Section IV.

It is known that Grohe's characterisation applies only to classes $\mathcal{C}$ of *bounded arity*, i.e., when the arities of the signatures are always bounded by a constant (for instance, CSPs over digraphs) and fails for classes of unbounded arity. In this direction, several hypergraph-based restrictions that lead to tractability have been proposed (for a survey see, e.g. [25]). Nevertheless, the precise tractability frontier for CSP($\mathcal{C}$, $-$) is not known. The situation is different for *fixed-parameter tractability*: Marx gave a complete classification of the fixed-parameter tractable restrictions CSP($\mathcal{C}$, $-$), for classes $\mathcal{C}$ of structures of unbounded arity [38]. In the case of VCSPs, Gottlob et al. [26] and Färnqvist [18] applied well-known hypergraph-based tractable restrictions of CSPs to VCSPs.

II. Preliminaries

We assume familiarity with relational structures and homomorphisms. Briefly, a *relational signature* is a finite set τ of relation symbols R, each with a specified arity $\text{ar}(R)$. A *relational structure* $\mathbf{A}$ over a relational signature τ (or a relational τ-structure, for short) is a finite universe A together with one relation $R^{\mathbf{A}} \subseteq A^{\text{ar}(R)}$ for each symbol $R \in \tau$. A *homomorphism* from a relational τ-structure $\mathbf{A}$ (with universe A) to a relational τ-structure $\mathbf{B}$ (with universe B) is a mapping $h : A \mapsto B$ such that for all $R \in \tau$ and all tuples $\mathbf{x} \in R^{\mathbf{A}}$ we have $h(\mathbf{x}) \in R^{\mathbf{B}}$. We refer the reader to [31] for more details.

We use $\overline{\mathbb{Q}}_{\geq 0}$ to denote the set of nonnegative rational numbers with positive infinity, i.e. $\overline{\mathbb{Q}}_{\geq 0} = \mathbb{Q}_{\geq 0} \cup \{\infty\}$. As usual, we assume that $\infty + c = c + \infty = \infty$ for all $c \in \overline{\mathbb{Q}}_{\geq 0}$, $\infty \times 0 = 0 \times \infty = 0$, and $\infty \times c = c \times \infty = \infty$, for all $c > 0$.

Valued structures: A *signature* is a finite set σ of function symbols f, each with a specified arity $\text{ar}(f)$. A *valued structure* $\mathbb{A}$ over a signature σ (or a valued σ-structure, for short) is a finite universe A together with one function $f^{\mathbb{A}} : A^{\text{ar}(f)} \mapsto \overline{\mathbb{Q}}_{\geq 0}$ for each symbol $f \in \sigma$. We define $\text{tup}(\mathbb{A})$ to be the set of all pairs $(f, \mathbf{x})$ such that $f \in \sigma$ and $\mathbf{x} \in A^{\text{ar}(f)}$. If $\mathbb{A}, \mathbb{B}, \ldots$ are valued structures, then $A, B, \ldots$ denote their respective universes.

VCSPs: We define *Valued Constraint Satisfaction Problems* (VCSPs) as in [44]. An instance of the VCSP is given by two valued structures $\mathbb{A}$ and $\mathbb{B}$ over the same signature σ. For a mapping $h : A \mapsto B$, we define

$$\text{cost}(h) = \sum_{(f,\mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(h(\mathbf{x})).$$

The goal is to find the minimum cost, denoted by opt$(\mathbb{A}, \mathbb{B})$, over all possible mappings $h : A \mapsto B$.

For a class $\mathcal{C}$ of valued structures (not necessarily over the same signature), we denote by VCSP($\mathcal{C}$, $-$) the restriction of the VCSP to instances $(\mathbb{A}, \mathbb{B})$ such that $\mathbb{A} \in \mathcal{C}$. We say that VCSP($\mathcal{C}$, $-$) is in PTIME, the class of problems solvable in *polynomial time*, if there is a deterministic algorithm that solves any instance $(\mathbb{A}, \mathbb{B})$ of VCSP($\mathcal{C}$, $-$) in time $(|\mathbb{A}| + |\mathbb{B}|)^{O(1)}$. We also consider the parameterised version of VCSP($\mathcal{C}$, $-$), denoted by p-VCSP($\mathcal{C}$, $-$), where the parameter is $|\mathbb{A}|$. We say that p-VCSP($\mathcal{C}$, $-$) is in FPT, the class of problems that are *fixed-parameter tractable*, if there is a deterministic algorithm that solves any instance $(\mathbb{A}, \mathbb{B})$ of p-VCSP($\mathcal{C}$, $-$) in time $f(|\mathbb{A}|) \cdot |\mathbb{B}|^{O(1)}$, where $f : \mathbb{N} \mapsto \mathbb{N}$ is an arbitrary computable function. The class W[1], introduced in [17], can be seen as an analogue of NP in parameterised complexity theory. Proving W[1]-hardness of p-VCSP($\mathcal{C}$, $-$) (under an fpt-reduction) is a strong indication that p-VCSP($\mathcal{C}$, $-$) is not in FPT as it is believed that FPT $\neq$ W[1]. We refer the reader to [22] for more details on parameterised complexity.

Treewidth of a valued structure: The notion of treewidth (originally introduced by Bertelé and Brioschi [4] and later rediscovered by Robertson and Seymour [40]) is a well-known measure of the tree-likeness of a graph [16]. Let $G = (V(G), E(G))$ be a graph. A *tree decomposition* of G is a pair (T, β) where $T = (V(T), E(T))$ is a tree and β is a function that maps each node $t \in V(T)$ to a subset of $V(G)$ such that (i) $V(G) = \bigcup_{t \in V(T)} \beta(t)$, (ii) for every $u \in V(G)$, the set $\{t \in V(T) \mid u \in \beta(t)\}$ induces a connected subgraph of T, and (iii) for every edge $\{u, v\} \in E(G)$, there is a node $t \in V(T)$ with $\{u, v\} \subseteq \beta(t)$. The *width* of the decomposition (T, β) is $\max\{|\beta(t)| \mid t \in V(T)\} - 1$. The *treewidth* $tw(G)$ of a graph G is the minimum width over all its tree decompositions.

Let $\mathbf{A}$ be a relational structure over a relational signature τ. Its *Gaifman graph* (also known as *primal graph*), denoted by $G(\mathbf{A})$, is the graph whose vertex set is the universe of $\mathbf{A}$ and whose edges are the pairs $\{u, v\}$ for which there is a tuple $\mathbf{x}$ and a relation symbol $R \in \tau$ such that u, v appear in $\mathbf{x}$ and $\mathbf{x} \in R^{\mathbf{A}}$. We define the treewidth of $\mathbf{A}$ to be $tw(\mathbf{A}) = tw(G(\mathbf{A}))$.

Let $\mathbb{A}$ be a valued σ-structure. We define the set of the *positive* tuples of $\mathbb{A}$ to be $\text{tup}(\mathbb{A})_{>0} = \{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}) \mid f^{\mathbb{A}}(\mathbf{x}) > 0\}$. Note that if $\mathbb{A}$ is the left-hand side of an instance of the VCSP, the only tuples relevant to the problem are those in $\text{tup}(\mathbb{A})_{>0}$. Hence, in order to define structural restrictions and in particular, the notion of treewidth, we focus on the structure induced by $\text{tup}(\mathbb{A})_{>0}$. Formally, we associate with the signature σ a relational signature rel(σ) that contains, for every $f \in \sigma$, a relation symbol R_f of the same arity as f. We define Pos$(\mathbb{A})$ to be the relational structure over rel(σ) with the same universe A of $\mathbb{A}$ such that $\mathbf{x} \in R_f^{\text{Pos}(\mathbb{A})}$ if and only if $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})_{>0}$. We let the treewidth of $\mathbb{A}$ be $tw(\mathbb{A}) = tw(\text{Pos}(\mathbb{A}))$.

Remark. Observe that, in the VCSP, we allow infinite costs not only in $\mathbb{B}$ but also in the left-hand side structure $\mathbb{A}$. This allows us to consider the VCSP as the minimum-cost mapping problem between two mathematical objects of the same nature. Intuitively, mapping the tuples of $\mathbb{A}$ to infinity ensures that those are logically equivalent to *hard constraints*, as any minimum-cost solution must map them to tuples of cost exactly 0 in $\mathbb{B}$. Thus, decision CSPs, which are $\{0, \infty\}$-valued VCSPs, are a special case of our definition and all our results also apply to CSPs.

III. EQUIVALENCE FOR VALUED STRUCTURES

We start by introducing the notion of valued equivalence that is crucial for our results.

Definition 1. *Let $\mathbb{A}, \mathbb{B}$ be valued σ-structures. We say that* $\mathbb{A}$ improves $\mathbb{B}$*, denoted by $\mathbb{A} \preceq \mathbb{B}$, if opt$(\mathbb{A}, \mathbb{C}) \leq$ opt$(\mathbb{B}, \mathbb{C})$ for all valued σ-structures $\mathbb{C}$.*

When two valued structures improve each other, we call them equivalent. (In Section I, we used the term "valued equivalence". From now on, we drop the word "valued" unless needed for clarity.)

Definition 2. *Let $\mathbb{A}, \mathbb{B}$ be valued σ-structures. We say that $\mathbb{A}$ and $\mathbb{B}$ are* equivalent*, denoted by $\mathbb{A} \equiv \mathbb{B}$, if $\mathbb{A} \preceq \mathbb{B}$ and $\mathbb{B} \preceq \mathbb{A}$.*

Hence, two valued σ-structures $\mathbb{A}$ and $\mathbb{B}$ are equivalent if they have the same optimal cost over all right-hand side valued structures. While equivalence implies *homomorphic* equivalence of Pos$(\mathbb{A})$ and Pos$(\mathbb{B})$, the converse does not hold in general [11] (cf. Example 2).

We now give a characterisation of equivalence in terms of certain types of homomorphisms. A homomorphism between two relational structures is a structure-preserving mapping. A *fractional* homomorphism between two valued structures played an important role in [44]. Intuitively, it is a probability distribution over mappings between the universes of the two structures with the property that the expected cost is not increased [44]. In this paper, we will need a different but related notion of *inverse* fractional homomorphism. For sets A and B, we denote by B^A the set of all mappings from A to B.

Definition 3. *Let $\mathbb{A}, \mathbb{B}$ be valued σ-structures. An* inverse fractional homomorphism *from $\mathbb{A}$ to $\mathbb{B}$ is a function ω :*

$B^A \mapsto \mathbb{Q}_{\geq 0}$ *with* $\sum_{g\in B^A}\omega(g) = 1$ *such that for each* $(f, \mathbf{x}) \in tup(\mathbb{B})$ *we have*

$$\sum_{g\in B^A} \omega(g) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) \leq f^{\mathbb{B}}(\mathbf{x})$$

where $f^{\mathbb{A}}(g^{-1}(\mathbf{x})) := \sum_{\mathbf{y}\in A^{ar(f)}:g(\mathbf{y})=\mathbf{x}} f^{\mathbb{A}}(\mathbf{y})$. *We define the* support *of* ω *to be the set* $supp(\omega) := \{g \in B^A \mid \omega(g) > 0\}$.

Observe that an inverse fractional homomorphism ω from $\mathbb{A}$ to $\mathbb{B}$ is actually a distribution over the set of homomorphisms from Pos($\mathbb{A}$) to Pos($\mathbb{B}$), i.e., every $g \in$ supp(ω) is a homomorphism from Pos($\mathbb{A}$) to Pos($\mathbb{B}$). The following result relates improvement and inverse fractional homomorphisms.

Proposition III.1. *Let* $\mathbb{A}, \mathbb{B}$ *be valued* σ*-structures. Then,* $\mathbb{A} \preceq \mathbb{B}$ *if and only if there exists an inverse fractional homomorphism from* $\mathbb{A}$ *to* $\mathbb{B}$.

Proof sketch: From right to left, take an arbitrary valued σ-structure $\mathbb{C}$ and a minimum-cost mapping h from $\mathbb{B}$ to $\mathbb{C}$. The idea is to compose the inverse fractional homomorphism ω from $\mathbb{A}$ to $\mathbb{B}$ with the mapping h to obtain a distribution over mappings from $\mathbb{A}$ to $\mathbb{C}$ whose cost in expectation is less than cost(h) = opt($\mathbb{B}, \mathbb{C}$). This implies that opt($\mathbb{A}, \mathbb{C}$) $\leq$ opt($\mathbb{B}, \mathbb{C}$). For the other direction, we can write a linear system capturing the existence of an inverse fractional homomorphism from $\mathbb{A}$ to $\mathbb{B}$. If there is no such fractional homomorphism, by Farkas' Lemma, we obtain a vector w of weights (one weight for each element in tup($\mathbb{B}$)) with certain properties. From $\mathbb{B}$ and w, we can define a counterexample to the fact that $\mathbb{A} \preceq \mathbb{B}$.

Appropriate notions of cores have played an important role in the complexity classifications of left-hand side restricted CSPs [29], right-hand side restricted CSPs [9], [7], [47], and right-hand side restricted VCSPs [45], [35]. We define cores around inverse fractional homomorphisms. We say that a valued structure $\mathbb{A}$ is a *core* if every inverse fractional homomorphism ω from $\mathbb{A}$ to $\mathbb{A}$ is *surjective*, i.e., every $g \in$ supp(ω) is surjective. A valued structure $\mathbb{A}'$ is a *core* of $\mathbb{A}$ if $\mathbb{A}'$ is a core itself and $\mathbb{A}' \equiv \mathbb{A}$. In the full version [11], we show that

(i) every valued structure $\mathbb{A}$ has a core,
(ii) all cores of $\mathbb{A}$ are isomorphic (in the sense of [11, Definition 8]),
(iii) *the* core of $\mathbb{A}$ is a structure with minimum number of elements among those equivalent to $\mathbb{A}$,
(iv) the core of $\mathbb{A}$ is a structure with minimum treewidth among those equivalent to $\mathbb{A}$, and
(v) the core of $\mathbb{A}$ can be computed effectively.

IV. Complexity of VCSP($\mathcal{C}$, $-$)

Let $\mathcal{C}$ be a class of valued structures. We say that $\mathcal{C}$ has *bounded arity* if there is a constant $r \geq 1$ such that for every valued σ-structure $\mathbb{A} \in \mathcal{C}$ and $f \in \sigma$, we have that ar(f) $\leq r$. Similarly, we say that $\mathcal{C}$ has bounded treewidth *modulo equivalence* if there is a constant $k \geq 1$ such that every $\mathbb{A} \in \mathcal{C}$ is equivalent to a valued structure $\mathbb{A}'$ with $tw(\mathbb{A}') \leq k$. The following is our first main result.

Theorem IV.1 (Complexity classification). *Assume FPT* $\neq$ *W[1]. Let* $\mathcal{C}$ *be a recursively enumerable class of valued structures of bounded arity. Then, the following are equivalent:*

1) VCSP*($\mathcal{C}$, $-$) is in PTIME.*
2) *p*-VCSP*($\mathcal{C}$, $-$) is in FPT.*
3) *$\mathcal{C}$ has bounded treewidth modulo equivalence.*

Although Grohe's result [29] for CSPs looks almost identical to Theorem IV.1, we emphasise that his result involves a different type of equivalence. In Grohe's case, the equivalence in question is *homomorphic equivalence* whereas in our case the equivalence in question involves *improvement*. As we will explain later in this section, Grohe's classification follows as a special case of Theorem IV.1.

Note that by property (iv) of cores discussed at the end of Section III, a class $\mathcal{C}$ has bounded treewidth modulo equivalence if and only if the class given by the *cores* of the valued structures in $\mathcal{C}$ has bounded treewidth. This notion strictly generalises bounded treewidth, as illustrated in Example 1. Consequently, Theorem IV.1 gives new tractable cases.

Example 1. Consider the signature $\sigma = \{f, \mu\}$, where f and μ are binary and unary function symbols, respectively. For $n \geq 1$, let $\mathbb{A}_n$ be the valued σ-structure with universe $A_n = \{1, \ldots, n\} \times \{1, \ldots, n\}$ such that (i) $f^{\mathbb{A}_n}((i,j),(i',j')) = \infty$ if $i \leq i'$, $j \leq j'$, and $(i'-i)+(j'-j) = 1$; otherwise $f^{\mathbb{A}_n}((i,j),(i',j')) = 0$, and (ii) $\mu^{\mathbb{A}_n}((i,j)) = 1$, for all $(i,j) \in A_n$. Also, for $n \geq 1$, let $\mathbb{A}'_n$ be the valued σ-structure with universe $A'_n = \{1, \ldots, 2n-1\}$ such that (i) $f^{\mathbb{A}'_n}(i,j) = \infty$ if $j = i+1$; otherwise $f^{\mathbb{A}'_n}(i,j) = 0$, and (ii) $\mu^{\mathbb{A}'_n}(i) = i$, for $1 \leq i \leq n$, and $\mu^{\mathbb{A}'_n}(i) = 2n-i$, for $n+1 \leq i \leq 2n-1$. The structures $\mathbb{A}$ and $\mathbb{A}'$ from Figure 1 correspond to $\mathbb{A}_3$ and $\mathbb{A}'_3$, respectively; informally $\mathbb{A}_n$ is a directed grid of size $n \times n$ with a unary function μ with weight 1 applied to each element. We argue in [11] that for each $n \geq 1$ the valued structure $\mathbb{A}'_n$ is the core of $\mathbb{A}_n$. Since $tw(\mathbb{A}'_n) = 1$, the class $\mathcal{C} := \{\mathbb{A}_n \mid n \geq 1\}$ has bounded treewidth modulo equivalence. However, $\mathcal{C}$ has unbounded treewidth as the Gaifman graphs in $\{G(\text{Pos}(\mathbb{A}_n)) \mid n \geq 1\}$ correspond to the class of (undirected) *grids*, which is a well-known family of graphs with unbounded treewidth (see, e.g. [16]).

It is worth noticing that bounded treewidth modulo equivalence implies bounded treewidth modulo homomorphic equivalence (of the positive parts), but the converse is not true in general, as the next example shows. Therefore, Theorem IV.1 tells us that the tractability frontier for VCSP($\mathcal{C}$,

$-$) lies *strictly* between bounded treewidth and bounded treewidth modulo homomorphic equivalence.

Example 2. For $n \geq 3$, let $\mathbb{A}_n$ be the valued σ-structure from Example 1. Let $\mathbb{C}_n$ be the valued σ-structure with the same universe as $\mathbb{A}_n$, i.e., $C_n = \{1,\dots,n\} \times \{1,\dots,n\}$, such that $f^{\mathbb{C}_n} = f^{\mathbb{A}_n}$ and $\mu^{\mathbb{C}_n}$ is defined as follows. Let $D_1,\dots,D_n$ be the n first diagonals of $\mathbb{C}_n$ starting from the bottom left corner $(1,1)$ (see Figure 1 for an illustration of $\mathbb{C}_4$). For $1 \leq i \leq n$, let E_i be the top-left to bottom-right enumeration of D_i. In particular, $E_1 = ((1,1))$, $E_2 = ((2,1),(1,2))$, $E_3 = ((3,1),(2,2),(1,3))$ and $E_n = ((n,1),(n-1,2),\dots,(1,n))$. Fix an integer $M = M(n)$ such that $M > n^2$. The values assigned by $\mu^{\mathbb{C}_n}$ to E_1, E_2 and E_3 are (1), $(M,1)$ and $\big(M^3, M^2, M^4\big)$, respectively, and for E_i, with $4 \leq i \leq n$, is $(M,1,\dots,1,M)$. All remaining elements in $C_n \setminus \bigcup_{1 \leq i \leq n} D_i$ receive cost 1. Figure 1 depicts the case of $\mathbb{C}_4$. Let $\mathcal{C} := \{\mathbb{C}_n \mid n \geq 3\}$. Note first that $\text{Pos}(\mathbb{C}_n)$ is homomorphically equivalent to the relational structure $\mathbf{P}_{2n-1}$ over relational signature $\text{rel}(\sigma) = \{R_f, R_\mu\}$ (recall the definition of $\text{rel}(\sigma)$ from Section II), whose universe is $P_{2n-1} = \{1,\dots,2n-1\}$, $R_\mu^{\mathbf{P}_{2n-1}} = P_{2n-1}$ and $R_f^{\mathbf{P}_{2n-1}} = \{(i,i{+}1) \mid 1 \leq i \leq 2n{-}2\}$. Since $tw(\mathbf{P}_{2n-1}) = 1$, for all $n \geq 3$, it follows that $\{\text{Pos}(\mathbb{C}_n) \mid n \geq 3\}$ has bounded treewidth modulo homomorphic equivalence. We show in [11] that $\mathbb{C}_n$ is a core, for all $n \geq 3$, and thus $\mathcal{C}$ has unbounded treewidth modulo (valued) equivalence.

Corollaries of the complexity classification: We can obtain the classification for CSPs of Dalmau et al. [14] and Grohe [29] as a special case of Theorem IV.1. Indeed, we can encode left-hand side relational structures as $\{0,\infty\}$-valued structures in the following way: a tuple has value ∞ if it belongs to the corresponding relation, otherwise its value is 0. Also, for a class $\mathcal{C}$ of $\{0,\infty\}$-valued structures, the problem VCSP($\mathcal{C}$,$-$) boils down to CSP($\mathcal{C}$,$-$), and the notion of bounded treewidth modulo (valued) equivalence collapses to bounded treewidth modulo homomorphic equivalence. Hence, by applying Theorem IV.1 to $\{0,\infty\}$-valued structures, we recover the known CSP classification from [14], [29].

In his PhD thesis [19], Färnqvist also considered the complexity of VCSP($\mathcal{C}$, $-$). However, he considered a different definition of the problem, that we denote by VCSP$_F$($\mathcal{C}$, $-$). In his framework, VCSPs are parameterised by a class of relational structures $\mathcal{C}$, instead of valued structures. In particular, VCSP$_F$($\mathcal{C}$, $-$) coincides with VCSP($\mathcal{C}_F$, $-$), where each valued structure $\mathbb{A} \in \mathcal{C}_F$ is obtained from a relational structure $\mathbf{A} \in \mathcal{C}$ by ignoring the signature of $\mathbf{A}$ and adding one fresh function symbol $f_{R,\mathbf{x}}$ for each tuple $\mathbf{x} \in R^{\mathbf{A}}$ such that $f^{\mathbb{A}}_{R,\mathbf{x}}(\mathbf{x}) = 1$ and $f^{\mathbb{A}}_{R,\mathbf{x}}(\mathbf{y}) = 0$, for all $\mathbf{y} \neq \mathbf{x}$. (For more details on this framework, see [11, Section 4].) It was shown in [19] that for a class $\mathcal{C}$ of relational structures of bounded arity, VCSP$_F$($\mathcal{C}$, $-$) is tractable if and only if $\mathcal{C}$ has bounded treewidth. This result follows directly from Theorem IV.1 as every valued structure in a class of the form $\mathcal{C}_F$ is a (valued) core, and hence, bounded treewidth modulo equivalence boils down to bounded treewidth. Notice that our definition of VCSP($\mathcal{C}$, $-$) is parameterised directly by a class of valued structures $\mathcal{C}$. This allows for a more fine-grained analysis of structural restrictions and, in particular, provides new tractable classes beyond bounded treewidth (cf. Example 1).

Finally, let us note that since Theorem IV.1 applies to *all* valued structures, the tractability part applies directly to the *finite-valued* VCSP, where all functions are restricted to take finite values in $\mathbb{Q}_{\geq 0}$. In fact, also the hardness part of Theorem IV.1 holds for the finite-valued VCSP [11]. Moreover, we give in [11] examples that demonstrate that already for finite-valued structures the tractability frontier is strictly between bounded treewidth and bounded treewidth modulo homomorphic equivalence.

Proof sketch of Theorem IV.1: The implication $(1) \Rightarrow (2)$ is immediate. The tractability part, i.e., implication $(3) \Rightarrow (1)$ will be established in Section V. In particular, it will follow from Theorem V.1 that, if there is a constant $k \geq 1$ such that every valued structure in the class $\mathcal{C}$ is equivalent to a valued structure of treewidth at most k, then VCSP($\mathcal{C}$, $-$) can be solved in polynomial time using the $(k+1)$-th level of the Sherali-Adams LP hierarchy.

Let us first mention that the remaining hardness part (implication $(2) \Rightarrow (3)$) does not follow directly from Grohe's result for CSPs [29]. The natural approach is to define, for a class of valued structures $\mathcal{C}$, the class of relational structures $\text{Pos}(\mathcal{C}) = \{\text{Pos}(\mathbb{A}) \mid \mathbb{A} \in \mathcal{C}\}$. Then one can observe that p-CSP(Pos($\mathcal{C}$), $-$) fpt-reduces to p-VCSP($\mathcal{C}$, $-$), and hence W[1]-hardness of the former problem implies hardness for the latter. However, if $\mathcal{C}$ has unbounded treewidth modulo equivalence, the class $\text{Pos}(\mathcal{C})$ does not necessarily have unbounded treewidth modulo homomorphic equivalence (see Example 2), and hence $\text{Pos}(\mathcal{C})$ is not necessarily hard according to Grohe's classification.

We instead refine Grohe's hardness proof [29] and reduce p-CLIQUE to p-VCSP($\mathcal{C}$, $-$), where $\mathcal{C}$ has unbounded treewidth modulo (valued) equivalence. Given an instance (G,k) of p-CLIQUE, the first step is to enumerate $\mathcal{C}$ until we find a valued structure $\mathbb{A}' \in \mathcal{C}$ with a (valued) core $\mathbb{A}$ such that the Gaifman graph $G(\mathbb{A})$ of $\mathbb{A}$ contains the $(k \times K)$-grid as a minor ($K = \binom{k}{2}$). Note that such an $\mathbb{A}$ exists due to the Excluded Grid Theorem [41] of Robertson and Seymour. Then we construct a valued structure $\mathbb{B}$ and a threshold $M^* \geq 0$ such that G contains a k-clique if and only if $\text{opt}(\mathbb{A},\mathbb{B}) \leq M^*$ (or equivalently, $\text{opt}(\mathbb{A}',\mathbb{B}) \leq M^*$), which completes the reduction. To define $\mathbb{B}$, we exploit a construction of [29] that, given an instance (G,k) of p-CLIQUE and a relational core $\mathbf{A}$ whose Gaifman graph contains the $(k \times K)$-grid as a minor, produces a relational structure $\mathbf{B}$ such that G contains a k-clique if and only if there is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$. (Intuitively, the

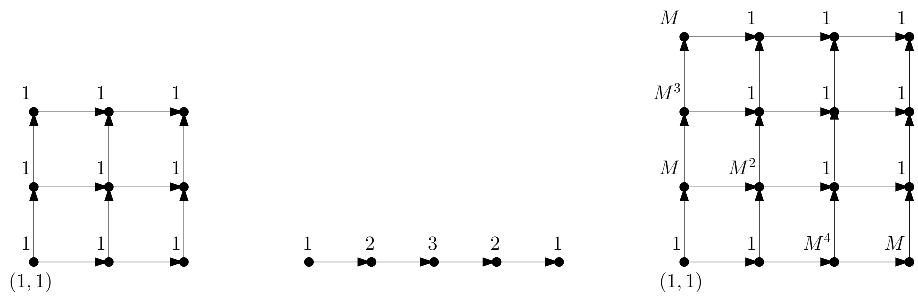

Figure 1. The valued σ-structures $\mathbb{A}_3$, $\mathbb{A}_3'$ and $\mathbb{C}_4$ from left to right ($M > 16$).

$(k \times K)$-grid minor of $\mathbf{A}$ encodes the incidence matrix of a k-clique.) Note that this construction cannot be applied directly with $\mathbf{A} = \text{Pos}(\mathbb{A})$ as being a (valued) core does not imply that the positive part is a relational core (see Example 2), and hence, $\text{Pos}(\mathbb{A})$ is not necessarily a relational core. Instead, we observe that it is possible to construct the above-described $\mathbf{B}$, together with a homomorphism π from $\mathbf{B}$ to $\text{Pos}(\mathbb{A})$, such that G contains a k-clique if and only if there is a homomorphism h from $\text{Pos}(\mathbb{A})$ to $\mathbf{B}$ such that $\pi \circ h$ is surjective [11, Lemma 25]. Finally, we exploit the following characterisation of cores [11, Proposition 15].

Proposition IV.2. *Let $\mathbb{A}$ be a valued σ-structure. Then, $\mathbb{A}$ is a core if and only if there exists a mapping c : $tup(\mathbb{A}) \mapsto \mathbb{Q}_{\geq 0}$ such that for every non-surjective mapping $g : A \mapsto A$, we have $M_{\mathbb{A},c} := \sum_{(f,\mathbf{x})\in tup(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x})c(f,\mathbf{x}) < \sum_{(f,\mathbf{x})\in tup(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x})c(f,g(\mathbf{x})) =: cost_{\mathbb{A},c}(g)$.*

Proposition IV.2 allows us to define our required threshold $M^* := M_{\mathbb{A},c}$ and the valued structure $\mathbb{B}$ by assigning values to the tuples of $\mathbf{B}$ according to c (via π), and large values to the remaining tuples. The key property of $\mathbb{B}$ is that for every assignment $h : A \mapsto B$, either h is not a homomorphism from $\text{Pos}(\mathbb{A})$ to $\mathbf{B}$, and then $\text{cost}(h)$ is large (in particular, larger than M^*), or h is a homomorphism and $\text{cost}(h) = \text{cost}_{\mathbb{A},c}(\pi \circ h)$. Using this, it can be shown that there is a homomorphism h from $\text{Pos}(\mathbb{A})$ to $\mathbf{B}$ such that $\pi \circ h$ is surjective if and only if $\text{opt}(\mathbb{A},\mathbb{B}) \leq M^*$.

V. Power of Sherali-Adams for VCSP($\mathcal{C}$, $-$)

Given a tuple $\mathbf{x}$, we write $\text{Set}(\mathbf{x})$ to denote the set of elements appearing in $\mathbf{x}$. Let $(\mathbb{A},\mathbb{B})$ be an instance of the VCSP over a signature σ and $k \geq 1$. We define a new signature $\sigma_k = \sigma \cup \{\rho_k\}$, where ρ_k is a new function symbol of arity k. Then, we create from $(\mathbb{A},\mathbb{B})$ an instance $(\mathbb{A}_k,\mathbb{B}_k)$ over σ_k such that $A_k = A$, $B_k = B$, $\rho_k^{\mathbb{A}_k}(\mathbf{x}) = 1$ for any $\mathbf{x} \in A_k^k$, $\rho_k^{\mathbb{B}_k}(\mathbf{x}) = 0$ for any $\mathbf{x} \in B_k^k$, and for every $f \in \sigma$ we have $f^{\mathbb{A}_k} = f^{\mathbb{A}}$ and $f^{\mathbb{B}_k} = f^{\mathbb{B}}$. Because the new function ρ_k is identically zero in $\mathbb{B}_k$, we have that for any mapping $h : A \mapsto B$, $\text{cost}(h)$ is the same in both instances $(\mathbb{A},\mathbb{B})$ and $(\mathbb{A}_k,\mathbb{B}_k)$. The *Sherali-Adams relaxation of level k* [43] of $(\mathbb{A},\mathbb{B})$, denoted by $\text{SA}_k(\mathbb{A},\mathbb{B})$, is the linear program given in Figure 2, which has one variable $\lambda(f,\mathbf{x},s)$ for each $(f,\mathbf{x}) \in \text{tup}(\mathbb{A}_k)_{>0}$ and $s : \text{Set}(\mathbf{x}) \mapsto B_k$. Note that the variables are indexed not only by $\mathbf{x}$ and s but also by f. This would not be necessary if $k \geq \text{ar}(f)$ but we are also interested in the case of $k < \text{ar}(f)$.

Definition 4. *Let $\mathbb{A},\mathbb{B}$ be valued σ-structures and $k \geq 1$. We denote by $opt_k(\mathbb{A},\mathbb{B})$ the minimum cost of a solution to $SA_k(\mathbb{A},\mathbb{B})$.*

Let $\mathbf{A}$ be a relational structure with universe A over a relational signature τ. Recall from Section II that $G(\mathbf{A})$ denotes the Gaifman graph of $\mathbf{A}$. A *scope* of $G(\mathbf{A})$ is a set X for which there is a relation symbol $R \in \tau$ and a tuple $\mathbf{x} \in R^{\mathbf{A}}$ such that $X = \text{Set}(\mathbf{x})$. In other words, the scopes of $G(\mathbf{A})$ are the sets that appear precisely in the tuples of $\mathbf{A}$.[1] Observe that every scope X of $G(\mathbf{A})$ induces a clique in $G(\mathbf{A})$.

Definition 5. *Let $\mathbf{A}$ be a relational structure and $G(\mathbf{A})$ its Gaifman graph. Let (T,β) be a tree decomposition of $G(\mathbf{A})$, where $T = (V(T),E(T))$. The* width modulo scopes *of (T,β) is defined as*

$$\max\{|\beta(t)| - 1 \mid t \in V(T) \text{ and } \beta(t) \text{ is not a scope of } G(\mathbf{A})\}.$$

If $\beta(t)$ is a scope for all nodes $t \in V(T)$ then we set the width modulo scopes of (T,β) to be 0. The treewidth modulo scopes *of $G(\mathbf{A})$, denoted by $tw_{ms}(G(\mathbf{A}))$, is the minimum width modulo scopes over all its tree decompositions. The treewidth modulo scopes of $\mathbf{A}$ is $tw_{ms}(\mathbf{A}) = tw_{ms}(G(\mathbf{A}))$. For a valued structure $\mathbb{A}$, we define the treewidth modulo scopes of $\mathbb{A}$ as $tw_{ms}(\mathbb{A}) = tw_{ms}(Pos(\mathbb{A}))$.*

Note that, unlike treewidth, the notion of treewidth modulo scopes is not *monotone*, i.e., it can increase after taking substructures. To see this, take for instance the relational structure $\mathbf{A}$ that corresponds to the undirected $k \times k$ grid. We have $tw_{ms}(\mathbf{A}) = k$. However, adding a new relation with only one tuple containing all elements of $\mathbf{A}$ lowers

[1]In a (V)CSP instance, the term *scope* usually refers to the list of variables a (valued) constraint depends on.

$$
\begin{aligned}
&\min \sum_{\substack{(f,\mathbf{x})\in \mathrm{tup}(\mathbb{A}_k)_{>0},\ s:\mathrm{Set}(\mathbf{x})\mapsto B_k \\ f^{\mathbb{A}_k}(\mathbf{x})\times f^{\mathbb{B}_k}(s(\mathbf{x}))<\infty}} \lambda(f,\mathbf{x},s) f^{\mathbb{A}_k}(\mathbf{x}) f^{\mathbb{B}_k}(s(\mathbf{x})) \\
&\lambda(f,\mathbf{x},s) = \sum_{r:\mathrm{Set}(\mathbf{y})\mapsto B_k, r|_{\mathrm{Set}(\mathbf{x})}=s} \lambda(p,\mathbf{y},r) && \forall (f,\mathbf{x}),(p,\mathbf{y}) \in \mathrm{tup}(\mathbb{A}_k)_{>0}: \\
& && \mathrm{Set}(\mathbf{x}) \subseteq \mathrm{Set}(\mathbf{y}) \text{ and } |\mathrm{Set}(\mathbf{x})| \leq k; \\
& && \forall s : \mathrm{Set}(\mathbf{x}) \mapsto B_k \\
&\sum_{s:\mathrm{Set}(\mathbf{x})\mapsto B_k} \lambda(f,\mathbf{x},s) = 1 && \forall (f,\mathbf{x}) \in \mathrm{tup}(\mathbb{A}_k)_{>0} \\
&\lambda(f,\mathbf{x},s) = 0 && \forall (f,\mathbf{x}) \in \mathrm{tup}(\mathbb{A}_k)_{>0}, s : \mathrm{Set}(\mathbf{x}) \mapsto B_k : f^{\mathbb{A}_k}(\mathbf{x}) \times f^{\mathbb{B}_k}(s(\mathbf{x})) = \infty \\
&\lambda(f,\mathbf{x},s) \geq 0 && \forall (f,\mathbf{x}) \in \mathrm{tup}(\mathbb{A}_k)_{>0}, s : \mathrm{Set}(\mathbf{x}) \mapsto B_k
\end{aligned}
$$

Figure 2. The Sherali-Adams relaxation of level k of $(\mathbb{A}, \mathbb{B})$.

the treewidth modulo scopes to 0. Let us also remark that the relational structures with treewidth modulo scopes 0 are precisely the relational structures whose underlying hypergraphs are α-acyclic (see e.g. [25]).

Given a valued σ-structure $\mathbb{A}$, the *overlap* of $\mathbb{A}$ is the largest integer m such that there exist $(f,\mathbf{x}),(p,\mathbf{y}) \in \mathrm{tup}(\mathbb{A})_{>0}$ with $(f,\mathbf{x}) \neq (p,\mathbf{y})$ and $|\mathrm{Set}(\mathbf{x}) \cap \mathrm{Set}(\mathbf{y})| = m$.

The following is our second main result.

Theorem V.1 (Power of Sherali-Adams). *Let $\mathbb{A}$ be a valued σ-structure and let $k \geq 1$. Let $\mathbb{A}'$ be the core of $\mathbb{A}$. Then, $opt_k(\mathbb{A},\mathbb{B}) = opt(\mathbb{A},\mathbb{B})$ for every valued σ-structure $\mathbb{B}$* if and only if *(i) $tw_{ms}(\mathbb{A}') \leq k-1$ and (ii) the overlap of $\mathbb{A}'$ is at most k.*

We show in [11] that Theorem V.1 holds also for finite-valued structures. In particular, the sufficiency part of Theorem V.1, i.e., if the core satisfies conditions (i) and (ii) then the k-th level of Sherali-Adams is tight, applies directly to the finite-valued case. It follows from the proofs of [11, Theorems 36 and Theorem 39] that whenever the core violates condition (i) or (ii) then the k-th level of Sherali-Adams is not tight even for finite-valued structures. Hence, Theorem V.1 also characterises the tightness of Sherali-Adams for finite-valued VCSPs.

Let us note that the characterisation given by Theorem V.1 for levels $k \geq r$, where r is the arity of the signature of $\mathbb{A}$, boils down to the notion of treewidth. That is, if $k \geq r$, the k-th level of Sherali-Adams is tight if and only if the treewidth of the core of $\mathbb{A}$ is at most $k-1$. Interestingly, Theorem V.1 tells us *precisely* under which conditions the k-th level works even for $k < r$.

Finally, we remark that the core structure $\mathbb{A}'$ is in general hard to compute [11], which rules out the naive algorithm that would compute $\mathrm{opt}(\mathbb{A},\mathbb{B})$ using dynamic programming along a tree decomposition of $\mathrm{Pos}(\mathbb{A}')$. Theorem V.1 gives a way to circumvent this issue since the linear program $\mathrm{SA}_k(\mathbb{A},\mathbb{B})$ does not depend on $\mathbb{A}'$ in any way.

Proof sketch of Theorem V.1: Suppose that $\mathbb{A}'$ is the core of $\mathbb{A}$ and that $\mathbb{A}'$ satisfies conditions (i) and (ii). Let (T,β) be a tree decomposition of $G(\mathrm{Pos}(\mathbb{A}'))$ of width modulo scopes at most $k-1$. Let λ be an optimal solution to $\mathrm{opt}_k(\mathbb{A}',\mathbb{B})$, for some arbitrary $\mathbb{B}$, and c be the vector defining the objective function of $\mathrm{SA}_k(\mathbb{A}',\mathbb{B})$. To show that $\mathrm{opt}(\mathbb{A}',\mathbb{B}) \leq \mathrm{opt}_k(\mathbb{A}',\mathbb{B})$, we consider the restriction of $\mathrm{SA}_k(\mathbb{A}',\mathbb{B})$ relevant to the tree decomposition (T,β). Formally, let $\mathrm{BTup} := \{(f,\mathbf{x}) \in \mathrm{tup}(\mathbb{A}'_k)_{>0} \mid \mathrm{Set}(\mathbf{x}) \subseteq \beta(t) \text{ for some } t \text{ in } T\}$ and $\mathcal{T} := \{(f,\mathbf{x},s) : (f,\mathbf{x}) \in \mathrm{BTup}, s : \mathrm{Set}(\mathbf{x}) \mapsto B\}$. Note that $\mathrm{tup}(\mathbb{A}')_{>0} \subseteq \mathrm{BTup}$, and hence $c|_{\mathcal{T}} \cdot \lambda|_{\mathcal{T}} = c \cdot \lambda = \mathrm{opt}_k(\mathbb{A}',\mathbb{B})$.

We show that the polytope described by the restriction of $\mathrm{SA}_k(\mathbb{A}',\mathbb{B})$ to $\mathcal{T}$ is integral (see [11, Lemma 34] for details). Note that a feasible solution of this polytope corresponds roughly to a family of consistent probability distributions $\mathcal{R} = \{R(X) \mid X \text{ is a bag of } (T,\beta)\}$, where each $R(X)$ is a distribution over partial mappings from A' to B with the same domain $X \subseteq A'$. Under this view, proving integrality (i.e., each feasible solution is a convex combination of integral solutions), corresponds to exhibiting a probability distribution $\mathcal{R}^*$ over (global) mappings from A' to B such that for every bag X of the tree decomposition (T,β), the marginal distribution of $\mathcal{R}^*$ over X coincides with $R(X)$. Since the overlap of $\mathbb{A}'$ is at most k, we have that $|X \cap X'| \leq k$ for every adjacent bags X and X' in (T,β). Together with the consistency constraints in $\mathrm{SA}_k(\mathbb{A}',\mathbb{B})$, this implies that the distributions $R(X)$ and $R(X')$ must be consistent (i.e., define the same marginal over $X \cap X'$). This allows us to argue inductively over the tree structure of T and construct the required distribution $\mathcal{R}^*$. This shows that $\mathrm{opt}(\mathbb{A}',\mathbb{B}) \leq \mathrm{opt}_k(\mathbb{A}',\mathbb{B})$ and hence $\mathrm{SA}_k(\mathbb{A}',\mathbb{B})$ is always tight. It follows that $\mathbb{A}$ is also tight as the optimum of any level of Sherali-Adams is preserved under equivalence [11, Proposition 28].

Next, we need to show that conditions (i) and (ii) are not only sufficient but also *necessary* for the tightness of the Sherali-Adams relaxation of level k. More precisely, given a valued structure $\mathbb{A}$ that fails to meet both conditions, we need to construct a valued structure $\mathbb{B}$ such that opt$(\mathbb{A},\mathbb{B}) >$ opt$_k(\mathbb{A},\mathbb{B})$. By [11, Proposition 28] we can further assume that $\mathbb{A}$ is a core.

In the case where $\mathbb{A}$ does not satisfy (i) we refine a proof strategy from [1], where it is shown that for every relational core $\mathbf{A}$ with treewidth at least k, there exists a relational structure $\mathbf{B}$ such that the $(k-1)$-consistency test on $(\mathbf{A},\mathbf{B})$ succeeds despite there being no homomorphism from $\mathbf{A}$ to $\mathbf{B}$. In order to prove this, it is shown that for every relational structure $\mathbf{A}$, there exists a relational structure $\mathbf{B}$ and a homomorphism π from $\mathbf{B}$ to $\mathbf{A}$, such that (†) there is no homomorphism h from $\mathbf{A}$ to $\mathbf{B}$ such that $\pi \circ h$ is surjective [1, Lemma 1]. Hence, if $\mathbf{A}$ is a relational core, there is no homomorphism from $\mathbf{A}$ to $\mathbf{B}$. On the other hand, if $tw(\mathbf{A}) \geq k$ then the $(k-1)$-consistency test on $(\mathbf{A},\mathbf{B})$ succeeds [1, Lemma 2], or equivalently, there is a non-empty family $\mathcal{H}_k$ of partial homomorphisms from $\mathbf{A}$ to $\mathbf{B}$ whose domains are of size at most k such that the following consistency property holds: for every two subsets $S, S' \subseteq A$ and $h \in \mathcal{H}_k$ with domain S, there is $h' \in \mathcal{H}_k$ with domain S' such that $h|_{S\cap S'} = h'|_{S\cap S'}$. To define such a family, the authors exploit the well-known connection between treewidth and brambles on graphs [42].

We define our valued structure $\mathbb{B}$ as a valued variant of the relational structure $\mathbf{B}$ that would be constructed from $\mathbf{A} = \text{Pos}(\mathbb{A})$, where the valuation of the tuples in $\mathbf{B}$ is given by the mapping c (via π) in our cost-based characterisation of cores (Proposition IV.2) and every other tuple is given a large value. The latter ensures that every mapping $h : A \mapsto B$ that is not a homomorphism from Pos$(\mathbb{A})$ to $\mathbf{B}$ has very large cost. On the other hand, every homomorphism h satisfies cost$(h) = \text{cost}_{\mathbb{A},c}(\pi \circ h)$. Property (†) above and Proposition IV.2 ensure that opt$(\mathbb{A},\mathbb{B}) > M_{\mathbb{A},c}$.

The delicate part of the proof is to argue that there exists a solution λ to SA$_k(\mathbb{A},\mathbb{B})$ of cost $M_{\mathbb{A},c}$. Let $\mathcal{D} := \{X \subseteq A \mid X \text{ is a scope of } G(\text{Pos}(\mathbb{A})) \text{ or } |X| \leq k\}$. The solution λ consists of a family of consistent probability distributions, each one defined over a set of partial mappings over the same domain X, where X ranges over $\mathcal{D}$. In order to define this family, we adapt an argument from [46] for proving the existence of gap instances for Sherali-Adams relaxations of VCSP$(-, E_{\mathcal{G},3})$, where $E_{\mathcal{G},3}$ denotes linear equations of width three. The idea is to define a well-behaved set of domains $\mathcal{S} \subseteq 2^A$ that covers $\mathcal{D}$ in the sense that for every $X \in \mathcal{D}$, there is a minimal set $S_X \in \mathcal{S}$ containing X. The set $\mathcal{S}$ is well-behaved in that one can define for $S \in \mathcal{S}$ a set of partial mappings $\mathcal{H}(S)$ with domain S, such that the family of uniform distributions $\{U(\mathcal{H}(S)) \mid S \in \mathcal{S}\}$ is consistent. This allows us to define our required family of consistent distributions over $\mathcal{D}$ via marginalisation: For each $X \in \mathcal{D}$, we consider the marginal distribution of $U(\mathcal{H}(S_X))$. The construction of $\{\mathcal{H}(S) \mid S \in \mathcal{S}\}$ is based on the family $\mathcal{H}_k$ from [1]. In our case, we exploit the fact that $tw_{ms}(\mathbb{A}) \geq k$ and a characterisation of treewidth modulo scopes in terms of brambles [11, Theorem 32]. Finally, to show that the uniform distributions $\{U(\mathcal{H}(S)) \mid S \in \mathcal{S}\}$ are actually consistent, we refine the analysis of $\mathcal{H}_k$ from [1]. In particular, it is not sufficient to argue that each partial mapping over domain S can be extended to one with domain S' but we need to reason about the number of such extensions [11, Claim 7].

If $\mathbb{A}$ violates condition (ii) instead, we construct a relational structure $\mathbf{B}$ over a domain B as follows. The elements of B are tuples of the form $(a, b_1, \ldots, b_n)$, where $a \in A$ and $(b_1, \ldots, b_n)$ is a list of bits. Then, we pick two distinct tuples $(p, \mathbf{x}), (q, \mathbf{y})$ with a large overlap $S := \text{Set}(\mathbf{x}) \cap \text{Set}(\mathbf{y})$ and define the relations of $\mathbf{B}$ with two properties in mind: For any homomorphism $h : \text{Pos}(\mathbb{A}) \mapsto \mathbf{B}$, the composition of h with the first projection is a homomorphism from Pos$(\mathbb{A})$ to Pos$(\mathbb{A})$, and there is no homomorphism $h : \text{Pos}(\mathbb{A}) \mapsto \mathbf{B}$ such that $|h(S)| = |S|$. The latter is achieved by exploiting $R_p^{\mathbf{B}}$ and $R_q^{\mathbf{B}}$ to enforce conflicting parity constraints on certain bits b_i in any injective assignment to S. Similarly to the case where condition (i) is violated, we then use these two properties together with the costs given by Proposition IV.2 to turn $\mathbf{B}$ into a valued structure $\mathbb{B}$ such that opt$(\mathbb{A},\mathbb{B})$ is strictly greater than some fixed threshold $M_{\mathbb{A},c}$. Finally, we exploit the fact that $R_p^{\mathbf{B}}$ and $R_q^{\mathbf{B}}$ only impose parity constraints on assignments to $|S| \geq k+1$ elements to show that SA$_k(\mathbb{A},\mathbb{B})$ has a solution of cost $M_{\mathbb{A},c}$ given by uniform distributions on carefully chosen sets of partial assignments.

VI. SEARCH VCSP($\mathcal{C}$, $-$)

If a class $\mathcal{C}$ of valued structures has bounded treewidth modulo equivalence then the Sherali-Adams LP hierarchy can be used to solve in polynomial time VCSP($\mathcal{C}$,$-$), that is, to compute the minimum cost of a mapping from $\mathbb{A} \in \mathcal{C}$ to some arbitrary valued structure $\mathbb{B}$. However, it may be the case that computing a mapping of that cost is NP-hard even though we know that one exists. The search version of the VCSP, denoted by SVCSP, explicitly asks for a minimum-cost mapping.

Given a valued σ-structure $\mathbb{A}$ and a mapping $g : A \mapsto A$, we define $g(\mathbb{A})$ to be the valued σ-structure over universe $g(A)$ such that $f^{g(\mathbb{A})}(\mathbf{x}) = f^{\mathbb{A}}(g^{-1}(\mathbf{x})) = \sum_{\mathbf{y} \in A^{\text{ar}(f)} : g(\mathbf{y}) = \mathbf{x}} f^{\mathbb{A}}(\mathbf{y})$, for all $f \in \sigma$ and $\mathbf{x} \in g(A)^{\text{ar}(f)}$. It follows from [11, Proposition 12] and [11, Proposition 42] that for every valued structure $\mathbb{A}$ there exists a mapping $g : A \mapsto A$ such that g belongs to the support of some inverse fractional homomorphism from $\mathbb{A}$ to $\mathbb{A}$ and $g(\mathbb{A})$ is the core of $\mathbb{A}$. If $\mathcal{C}$ is a class of valued structures, we denote by CORE COMPUTATION($\mathcal{C}$) the problem that takes as input some $\mathbb{A} \in \mathcal{C}$ and asks to compute such a mapping g.

Building on our results from Section V and adapting techniques from [45], we show the following.

Theorem VI.1 (**Search classification**). *Assume FPT $\neq$ W[1]. Let $\mathcal{C}$ be a recursively enumerable class of valued structures of bounded arity. Then, the following are equivalent:*

1) SVCSP($\mathcal{C}$, $-$) *is in PTIME.*
2) $\mathcal{C}$ *is of bounded treewidth modulo equivalence and* CORE COMPUTATION($\mathcal{C}$) *is in PTIME.*

Proof sketch: We start with the implication (2) $\Rightarrow$ (1). Suppose that $\mathcal{C}$ has treewidth modulo equivalence at most k for some $k \in \mathbb{N}$ and that CORE COMPUTATION($\mathcal{C}$) is in PTIME. Let $(\mathbb{A}, \mathbb{B})$ be an instance of SVCSP($\mathcal{C}$, $-$). First, we invoke the polynomial-time algorithm for CORE COMPUTATION($\mathcal{C}$) on the structure $\mathbb{A}$. As observed in [11, Proposition 10], composing the mapping g thus obtained with any minimum-cost mapping h from $g(\mathbb{A})$ to $\mathbb{B}$ yields a minimum-cost mapping from $\mathbb{A}$ to $\mathbb{B}$. We then combine a straightforward branching algorithm with Theorem V.1 to construct an integral optimal solution to $\mathrm{SA}_{k+1}(g(\mathbb{A}), \mathbb{B})$ in polynomial time, which provides a suitable mapping h. For this step, the key property is that introducing any set of fresh unary functions to $g(\mathbb{A})$ (for branching purposes) cannot increase its treewidth modulo equivalence because it is a core.

One part of the implication (1) $\Rightarrow$ (2) is immediate by Theorem IV.1: if SVCSP($\mathcal{C}$, $-$) is in PTIME, then $\mathcal{C}$ is of bounded treewidth modulo equivalence (under our assumptions). For the second part we proceed in two steps. First, we show that there is a polynomial-time Turing reduction from CORE COMPUTATION($\mathcal{C}$) to a slightly different problem REDUCTION STEP($\mathcal{C}$), which takes as input a structure $\mathbb{A} \in \mathcal{C}$ together with a mapping $g : A \mapsto A$ that belongs to the support of at least one inverse fractional homomorphism from $\mathbb{A}$ to $\mathbb{A}$. The goal of REDUCTION STEP is to compute a mapping g^+ that also belongs to the support of at least one inverse fractional homomorphism from $\mathbb{A}$ to $\mathbb{A}$ and such that $g^+(A) \subsetneq g(A)$. Adapting a strategy from [45], we reformulate the problem of computing such a mapping g^+ as a linear program (with exponentially many variables) and show that its dual admits a polynomial-time strong separation oracle. This oracle makes crucial use of the assumed polynomial-time algorithm for SVCSP($\mathcal{C}$, $-$). It then follows from standard combinatorial optimisation techniques [30, Lemma 6.5.15] that REDUCTION STEP($\mathcal{C}$) (and hence CORE COMPUTATION($\mathcal{C}$)) can be solved in polynomial time, which establishes the claim.

VII. RELATED PROBLEMS

In this section we provide tight complexity bounds for several problems related to our characterisations. (For complete proofs, we refer the reader to [11].) All lower bounds follow from existing results on relational structures or $\{0, \infty\}$-valued structures. For the upper bounds, we need the results from previous sections, especially, the machinery developed in Section III.

Proposition VII.1. *We have the following:*

1) *Given two valued structures $\mathbb{A}$, $\mathbb{B}$, deciding whether $\mathbb{A}$ improves $\mathbb{B}$ is NP-complete. Similarly, deciding whether $\mathbb{A}$ and $\mathbb{B}$ are equivalent is NP-complete.*
2) *Given a valued structure $\mathbb{A}$, deciding whether $\mathbb{A}$ is a core is coNP-complete.*
3) *For every fixed $k \geq 1$, the following problems are NP-complete:*
   a) *Given a valued structure $\mathbb{A}$, decide whether the treewidth of the core of $\mathbb{A}$ is at most k.*
   b) *Given a valued structure $\mathbb{A}$, decide whether the Sherali-Adams relaxation of level k is always tight for $\mathbb{A}$.*

VIII. APPLICATION TO DATABASE THEORY

It is well known that the evaluation/containment problem for *conjunctive queries* (CQs) (i.e., first-order queries using only conjunction and existential quantification) is equivalent to the homomorphism problem, and hence equivalent to CSPs [12], [33]. This observation has been fundamental in providing principled techniques for the static analysis and optimisation of CQs. Indeed, in their seminal work [12], Chandra and Merlin exploited this connection to show that the containment and equivalence problem for CQs are NP-complete. They also provided tools for minimising CQs with strong theoretical guarantees. In terms of homomorphisms, minimising a CQ corresponds essentially to computing the (relational) core of a relational structure.

The situation is less clear in the context of annotated databases [28]. In this framework, the tuples of the database are annotated with values from a particular semiring $\mathcal{K}$, and the semantics of a CQ is a value from $\mathcal{K}$. For instance, the Boolean semiring $(\{0, 1\}, \vee, \wedge, 0, 1)$ gives us the usual semantics of CQs, and the natural semiring $(\mathbb{N}, +, \times, 0, 1)$ corresponds to the so-called *bag semantics* of CQs. Another semiring considered in the literature is the *tropical* semiring $(\overline{\mathbb{Q}}_{\geq 0}, \min, +, \infty, 0)$, which provides a *minimum-cost* semantics [28]. Unfortunately, the homomorphism machinery cannot be applied directly to the study of containment and equivalence in the semiring setting. While there are some works in this direction (see, e.g. [36], [27]), several basic problems remain open. In particular, the precise complexity of containment/equivalence of CQs over the tropical semiring is open (it was shown in [36] to be NP-hard and in Π_2^p, the second level of the polynomial-time hierarchy). Our first observation is that these two problems are actually NP-complete. Indeed, it is well known that VCSP is equivalent to CQ evaluation over the tropical semiring. Moreover, containment and equivalence of CQs over the tropical semiring correspond to improvement and (valued) equivalence of valued structures. By applying Proposition VII.1, item (1), we directly obtain NP-completeness of these problems.

Our second observation is that our notion of (valued) core provides a notion of minimisation of CQs over the tropical semiring with theoretical guarantees. Indeed, as pointed out in item (iii) at the end of Section III, the core of a valued structure is always an equivalent valued structure with minimal number of elements, or in terms of CQs, with minimal number of variables. Also, as the core is computable (item (v) at the end of Section III), we have an algorithm to compute the core of a CQ over the tropical semiring. (In fact, a PSPACE algorithm.) Finally, it is worth mentioning that our classification result from Theorem IV.1 can be interpreted as a characterisation of the classes of CQs over the tropical semiring that can be evaluated in PTIME.

ACKNOWLEDGMENT

Stanislav Živný was supported by a Royal Society University Research Fellowship. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714532). The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

REFERENCES

[1] A. Atserias, A. A. Bulatov, and V. Dalmau, "On the Power of k-Consistency," in *Proceedings of the 34th International Colloquium on Automata, Languages and Programming (ICALP'07)*, ser. Lecture Notes in Computer Science, vol. 4596. Springer, 2007, pp. 279–290.

[2] L. Barto and M. Kozik, "Constraint Satisfaction Problems Solvable by Local Consistency Methods," *Journal of the ACM*, vol. 61, no. 1, 2014, article No. 3.

[3] L. Barto, M. Kozik, and T. Niven, "The CSP dichotomy holds for digraphs with no sources and no sinks (a positive answer to a conjecture of Bang-Jensen and Hell)," *SIAM Journal on Computing*, vol. 38, no. 5, pp. 1782–1802, 2009.

[4] U. Bertelé and F. Brioschi, *Nonserial dynamic programming*. Academic Press, 1972.

[5] D. Bienstock and N. Özbay, "Tree-width and the Sherali-Adams operator," *Discrete Optimization*, vol. 1, no. 1, pp. 13–21, 2004.

[6] A. Bulatov, "A dichotomy theorem for constraint satisfaction problems on a 3-element set," *Journal of the ACM*, vol. 53, no. 1, pp. 66–120, 2006.

[7] ——, "A dichotomy theorem for nonuniform CSPs," in *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*. IEEE, 2017, pp. 319–330.

[8] A. Bulatov, V. Dalmau, M. Grohe, and D. Marx, "Enumerating homomorphisms," *Journal of Computer and System Sciences*, vol. 78, no. 2, pp. 638–650, 2012.

[9] A. Bulatov, P. Jeavons, and A. Krokhin, "Classifying the Complexity of Constraints using Finite Algebras," *SIAM Journal on Computing*, vol. 34, no. 3, pp. 720–742, 2005.

[10] A. A. Bulatov, "Complexity of conservative constraint satisfaction problems," *ACM Transactions on Computational Logic*, vol. 12, no. 4, 2011, article 24.

[11] C. Carbonnel, M. Romero, and S. Živný, "The complexity of general-valued CSPs seen from the other side," Tech. Rep., October 2017, arXiv:1710.03148v1.

[12] A. K. Chandra and P. M. Merlin, "Optimal implementation of conjunctive queries in relational data bases," in *Proceedings of the 9th Annual ACM Symposium on Theory of Computing (STOC'77)*. ACM, 1977, pp. 77–90.

[13] D. A. Cohen, M. C. Cooper, P. G. Jeavons, and A. A. Krokhin, "The Complexity of Soft Constraint Satisfaction," *Artificial Intelligence*, vol. 170, no. 11, pp. 983–1016, 2006.

[14] V. Dalmau, P. G. Kolaitis, and M. Y. Vardi, "Constraint Satisfaction, Bounded Treewidth, and Finite-Variable Logics," in *Proceedings of the 8th International Conference on Principles and Practice of Constraint Programming (CP'02)*, ser. Lecture Notes in Computer Science, vol. 2470. Springer, 2002, pp. 310–326.

[15] R. Dechter, *Constraint Processing*. Morgan Kaufmann, 2003.

[16] R. Diestel, *Graph Theory*, 4th ed. Springer, 2010.

[17] R. G. Downey and M. R. Fellows, "Fixed-Parameter Tractability and Completeness I: Basic Results," *SIAM Journal on Computing Computing*, vol. 24, no. 4, pp. 873–921, 1995.

[18] T. Färnqvist, "Constraint optimization problems and bounded tree-width revisited," in *Proceedings of the 9th International Conference on Integration of Artificial Intelligence and Operations Research Techniques in Constraint Programming (CPAIOR'12)*, ser. Lecture Notes in Computer Science, vol. 7298. Springer, 2012, pp. 163–197.

[19] ——, "Exploiting Structure in CSP-related Problems," Ph.D. dissertation, Department of Computer Science and Information Science, Linköping University, 2013.

[20] T. Färnqvist and P. Jonsson, "Bounded tree-width and CSP-related problems," in *Proceedings of the 18th International Symposium on Algorithms and Computation (ISAAC)*, ser. Lecture Notes in Computer Science, vol. 4835. Springer, 2007, pp. 632–643.

[21] 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*, vol. 28, no. 1, pp. 57–104, 1998.

[22] J. Flum and M. Grohe, *Parameterized complexity theory*. Springer, 2006.

[23] E. C. Freuder, "A Sufficient Condition for Backtrack-free Search," *Journal of the ACM*, vol. 29, no. 1, pp. 24–32, 1982.

[24] ——, "Complexity of K-Tree Structured Constraint Satisfaction Problems," in *Proceedings of the 8th National Conference on Artificial Intelligence (AAAI'90)*, 1990, pp. 4–9.

[25] G. Gottlob, G. Greco, N. Leone, and F. Scarcello, "Hypertree decompositions: Questions and answers," in *Proceedings of the 35th SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'16)*, 2016, pp. 57–74.

[26] G. Gottlob, G. Greco, and F. Scarcello, "Tractable Optimization Problems through Hypergraph-Based Structural Restrictions," in *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP'09), Part II*, ser. Lecture Notes in Computer Science, vol. 5556. Springer, 2009, pp. 16–30.

[27] T. J. Green, "Containment of conjunctive queries on annotated relations," *Theory of Computing Systems*, vol. 49, no. 2, pp. 429–459, 2011.

[28] T. J. Green and V. Tannen, "The semiring framework for database provenance," in *Proceedings of the 36th SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'17)*, 2017, pp. 93–99.

[29] M. Grohe, "The complexity of homomorphism and constraint satisfaction problems seen from the other side," *Journal of the ACM*, vol. 54, no. 1, pp. 1–24, 2007.

[30] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics series*. Springer, 1988.

[31] P. Hell and J. Nešetřil, *Graphs and Homomorphisms*. Oxford University Press, 2004.

[32] P. M. Idziak, P. Markovic, R. McKenzie, M. Valeriote, and R. Willard, "Tractability and learnability arising from algebras with few subpowers," *SIAM Journal on Computing*, vol. 39, no. 7, pp. 3023–3037, 2010.

[33] P. G. Kolaitis and M. Y. Vardi, "Conjunctive-query containment and constraint satisfaction," in *Proceedings of the 17th SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'98)*, 1998, pp. 205–213.

[34] P. Kolman and M. Koutecký, "Extended formulation for CSP that is compact for instances of bounded treewidth," *Electr. J. Comb.*, vol. 22, no. 4, p. P4.30, 2015.

[35] V. Kolmogorov, A. A. Krokhin, and M. Rolínek, "The complexity of general-valued CSPs," *SIAM Journal on Computing*, vol. 46, no. 3, pp. 1087–1110, 2017.

[36] E. V. Kostylev, J. L. Reutter, and A. Z. Salamon, "Classification of annotation semirings over containment of conjunctive queries," *ACM Trans. Database Syst.*, vol. 39, no. 1, pp. 1:1–1:39, 2014.

[37] M. Kozik and J. Ochremiak, "Algebraic properties of valued constraint satisfaction problem," in *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP'15)*, ser. Lecture Notes in Computer Science, vol. 9134. Springer, 2015, pp. 846–858.

[38] D. Marx, "Tractable hypergraph properties for constraint satisfaction and conjunctive queries," *Journal of the ACM*, vol. 60, no. 6, 2013, article No. 42.

[39] U. Montanari, "Networks of Constraints: Fundamental properties and applications to picture processing," *Information Sciences*, vol. 7, pp. 95–132, 1974.

[40] N. Robertson and P. D. Seymour, "Graph minors. III. Planar tree-width," *Journal of Combinatorial Theory, Series B*, vol. 36, no. 1, pp. 49–64, 1984.

[41] ——, "Graph minors. V. Excluding a planar graph," *Journal of Combinatorial Theory, Series B*, vol. 41, no. 1, pp. 92–114, 1986.

[42] P. D. Seymour and R. Thomas, "Graph searching and a min-max theorem for tree-width," *Journal of Combinatorial Theory, Series B*, vol. 58, no. 1, pp. 22–33, 1993.

[43] H. D. Sherali and W. P. Adams, "A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems," *SIAM Journal of Discrete Mathematics*, vol. 3, no. 3, pp. 411–430, 1990.

[44] J. Thapper and S. Živný, "The power of linear programming for valued CSPs," in *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'12)*. IEEE, 2012, pp. 669–678.

[45] ——, "The complexity of finite-valued CSPs," *Journal of the ACM*, vol. 63, no. 4, 2016, article No. 37.

[46] ——, "The power of Sherali-Adams relaxations for general-valued CSPs," *SIAM Journal on Computing*, vol. 46, no. 4, pp. 1241–1279, 2017.

[47] D. Zhuk, "A proof of CSP dichotomy conjecture," in *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science (FOCS'17)*. IEEE, 2017, pp. 331–342.