

PTAS for Sparse General-Valued CSPs

BALÁZS F. MEZEI

MARCIN WROCHNA, University of Warsaw, Poland

STANISLAV ŽIVNÝ, University of Oxford, United Kingdom

We study polynomial-time approximation schemes (PTASes) for constraint satisfaction problems (CSPs) such as Maximum Independent Set or Minimum Vertex Cover on sparse graph classes.

Baker’s approach gives a PTAS on planar graphs, excluded-minor classes, and beyond. For Max-CSPs, and even more generally, maximisation finite-valued CSPs (where constraints are arbitrary non-negative functions), Romero, Wrochna, and Živný [SODA’21] showed that the Sherali-Adams LP relaxation gives a simple PTAS for all fractionally-treewidth-fragile classes, which is the most general “sparsity” condition for which a PTAS is known. We extend these results to general-valued CSPs, which include “crisp” (or “strict”) constraints that have to be satisfied by every feasible assignment. The only condition on the crisp constraints is that their domain contains an element which is at least as feasible as all the others (but possibly less valuable).

For minimisation general-valued CSPs with crisp constraints, we present a PTAS for all *Baker* graph classes – a definition by Dvořák [SODA’20] which encompasses all classes where Baker’s technique is known to work, except for fractionally-treewidth-fragile classes. While this is standard for problems satisfying a certain monotonicity condition on crisp constraints, we show this can be relaxed to *diagonalisability* – a property of relational structures connected to logics, statistical physics, and random CSPs.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms**.

Additional Key Words and Phrases: constraint satisfaction, Baker classes, sparsity, diagonalisability

ACM Reference Format:

Balázs F. Mezei, Marcin Wrochna, and Stanislav Živný. 2022. PTAS for Sparse General-Valued CSPs. *ACM Trans. Algor.* 1, 1, Article 1 (January 2022), 31 pages.

1 INTRODUCTION

Min-Ones and Max-Ones, studied by Khanna and Motwani (under the names of TMIN and TMAX, respectively) [31] and by Khanna, Sudan, Trevisan, and Williamson [32], are Boolean CSPs in which one seeks a feasible solution (a 0–1 assignment satisfying all constraints) minimising/maximising the number of variables assigned the label 1. Classical examples are the Minimum Vertex Cover and the Maximum Independent Set problem, respectively. A natural generalisation to larger alphabets is the problem in which one seeks a solution to a CSP instance while minimising/maximising a sum of unary functions. With injective unary functions, such problems have been studied under the name of Strict-CSP by plumar, Manokaran, Tulsiani, and Vishnoi [36], and Min/Max-Solution by Jonsson, Kuivinen, and Nordh [29]. With arbitrary unary functions, such problems have been studied under the name of Min-Cost-Hom by Gutin, Hell, Rafiey, and Yeo [23], Takhanov [46], and others [24, 25, 40]. In this paper we consider the still more general setting of *general-valued*

Authors’ addresses: Balázs F. Mezei; Marcin Wrochna, m.wrochna@mimuw.edu.pl, University of Warsaw, Warsaw, Poland; Stanislav Živný, standa.zivny@cs.ox.ac.uk, University of Oxford, Oxford, United Kingdom.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

1549-6325/2022/1-ART1 \$15.00

<https://doi.org/>

CSPs, where constraints are functions which give values to every possible assignment on a tuple of variables; we allow ∞ or $-\infty$ values to express *crisp* (also known as *strict*) constraints, which have to be satisfied by every feasible (finite-valued) assignment. While a lot of research is devoted to exact algorithms or optimal approximation ratios in APX-hard cases (see [27, 30, 37] for surveys), we seek the most general conditions that allow to obtain a polynomial-time approximation scheme (PTAS).

Baker [1] gave an elegant method (sometimes known as the *shifting* or *layering technique*) for constructing polynomial-time approximation schemes (PTASes) which applies to many such problems, with the condition that the input instance’s graph (the Gaifman graph) is “sparse”. This was initially presented for planar graphs, but it is known that similar structural properties are exhibited by all proper minor-closed graph classes [11, 12, 20] and beyond: e.g. graphs embeddable in a fixed surface with few intersections per edge [19, 38], or sparse unit ball intersection graphs in few dimensions [26] (but not e.g. 3-regular expanders: bounded degree is not sufficient to get a PTAS even for Independent Set [2]). Dvořák [15] defined *fractionally-treewidth-fragile* classes – a natural generalisation of earlier sparsity conditions – which encompasses all these examples. A class of graphs is *fractionally-treewidth-fragile* if one can remove vertices in a randomised way so that each vertex is removed with arbitrarily small probability ε , but the treewidth after removal is always bounded, the bound depending on ε only. He showed that if this notion of sparsity can be efficiently certified in a class of graphs, then this suffices to guarantee a PTAS, at least for a few problems such as Weighted Maximum Independent Set. On the other hand it is not known whether this suffices for Minimum Vertex Cover, for example.

To remedy this, Dvořák [16] later defined *Baker classes* and proved that (an effective version of) this condition suffices to provide a PTAS to all monotone optimisation problems expressible in first-order logic (including of course Vertex Cover). Very roughly, a class of graphs is *Baker* if one can reduce each graph in it to the empty graph by a bounded number of the following steps: either remove a single vertex, or select a breadth-first-search layering and recurse into all subgraphs that can be induced by a few consecutive layers. Dvořák proved that the family of Baker classes still includes all the examples discussed above; on the other hand, it is strictly included in the family of *fractionally-treewidth-fragile* classes (and hence less general) [14]. It is worth mentioning that proper minor-closed graph classes can be shown to be Baker (and *fractionally-treewidth-fragile*) relatively easily, without using the Graph Minor Structure theorem, in contrast to the earlier, less general definitions (see [16] for details).

In order to provide a PTAS for a class of general-valued CSPs, *a sparsity condition is not enough*: we also need to restrict what types of constraints can be used in an instance. Otherwise, even if the values to be optimised are trivial, either 0 or infinity, one could use the crisp constraints to express 3-Colouring, which is NP-hard even on planar graphs of bounded degree [10]. In fact, as long as all crisp constraints are available, for any possible restriction on Gaifman graphs, either the restriction implies bounded treewidth, making the problem exactly solvable, or it is hard to decide whether the optimum is zero or infinite, by a result of Grohe, Schwentick, and Segoufin [22]. We will hence require a condition which ensures that one can easily decide whether a feasible solution (of finite value) exists. This usually takes the form of a monotonicity condition.

On the other hand, *some sparsity condition is also necessary*: on general Gaifman graphs, there is no restriction of constraint types that would result in a general-valued CSP that admits a PTAS but is not solvable exactly in polynomial time.¹ In this sense our work follows the line of “uniform”

¹This follows from the NP-hardness result of Kozik and Ochremiak [35], which actually shows APX-hardness; for earlier, explicit APX-hardness results for CSPs see, e.g., [28, 29]. However, we remark non-trivial PTAS examples are known for “surjective” maximisation finite-valued CSPs [18].

or “hybrid” CSPs, which include restrictions on both the input’s Gaifman graph (*left-hand side* restrictions) and on the types of constraints (*right-hand side* restrictions); see [9] for a survey. However, unlike that line of work, we look for PTASes instead of exact solvability, which also lets us go well beyond planar graphs and beyond very specialised algebraic algorithms.

1.1 Related work

The exact solvability of general-valued CSPs has been characterised for left-hand side restrictions [8] (tractable cases are precisely classes that have bounded treewidth, up to a certain notion of homomorphic equivalence) and right-hand side restrictions [34] (tractable cases are precisely delineated by certain algebraic properties); both results include the case where infinite values are allowed.

As discussed above, there are no PTASes for general-valued CSPs with only left-hand side or only right-hand side restrictions, beyond exactly solvable cases. In fact Khanna et al. [32], in their work on Min-Ones and Max-Ones with right-hand side restriction, remark that “Our framework lacks such phenomena as PTAS” and discuss left-hand side restrictions as an interesting avenue for future work for that reason. Similarly [29] and [30] ask in the context of right-hand side restricted Min-Solution and Max-Solution problems: “Under which restrictions on variable scopes does Max Sol admit a PTAS?”

Very recently, PTASes for left-hand side restricted Max-CSP without crisp constraints, such as Max-Cut, have been studied by Romero, Wrochna, and Živný [44]. More generally, they consider so-called finite-valued CSPs, where the only right-hand side restriction is having finite, non-negative values. They showed a PTAS is possible for every fractionally-treewidth-fragile class of Gaifman graphs. In fact the algorithm is simply the Sherali-Adams linear programming relaxation (with a growing number of levels giving a better and better approximation), which is oblivious to the graph structure and does not require it to be efficiently certified in any way.

As for constant-factor approximations, Raghavendra’s celebrated result gave the best approximation ratio, assuming the Unique Games Conjecture of Khot [33], for all right-hand side restricted Max-CSPs (and also finite-valued CSPs) [41]. Analogous results for monotone Strict-CSPs were obtained by Kumar et al. [36]. Constant-factor approximation algorithms have been established for right-hand side restricted Min-Cost-Hom on special graphs by Hell, Mastrolilli, Nevisi, and Rafiey [24], and for all graphs and some digraphs by Rafiey, Rafiey, and Santos [40].

1.2 Our results

As in most research on constraint satisfaction, our main motivation is to understand the mathematical structure that governs efficient computation and, if at all possible, to obtain complexity classifications of large fragments of CSPs. In this paper, the goal was to push the tractability frontier of general-valued CSPs that admit polynomial-time approximation schemes. In particular, we try to understand what right-hand side restrictions make Baker’s technique possible, tentatively answering the aforementioned questions from [29, 30, 32]. We show that the most general results known for Vertex Cover and Independent Set (PTASes on all Baker or fractionally-treewidth-fragile classes, respectively) can be extended to any to general-valued CSPs with a certain monotonicity restriction, and even further.

To clearly separate left-hand side and right-hand side restriction, it is convenient to phrase a general-valued CSP (VCSP) as the problem of optimising the value of a function between two valued structures. Precise definitions are given in Section 2. Briefly, a *valued structure* \mathbb{A} consists of a (finite) domain A and a collection of functions $f^{\mathbb{A}}: A^n \rightarrow \mathbb{Q} \cup \{\pm\infty\}$, indexed by symbols f belonging to a set of symbols σ called a signature. For two (finite) structures \mathbb{A}, \mathbb{C} , the *value* of an assignment $h: A \rightarrow C$ is an expression of the form $\sum f^{\mathbb{A}}(\mathbf{x})f^{\mathbb{C}}(h(\mathbf{x}))$. We will be seeking

to find either the minimum or maximum value over all assignments, denoted $\text{minval}(\mathbb{A}, \mathbb{C})$ and $\text{maxval}(\mathbb{A}, \mathbb{C})$ respectively. *Feasible* assignments are those of finite value. The reader should think of the left-hand side structure \mathbb{A} as of a set of variables A together with weighted constraint scopes: for $\mathbf{x} \in A^n$, $f^{\mathbb{A}}(\mathbf{x}) = w \neq 0$ means that the instance applies the constraint “ f ” to variables in \mathbf{x} with weight w . The right-hand side structure \mathbb{C} encodes the alphabet C (to which an assignment h maps each variable) and the collection of available constraints, which could be arbitrary $\mathbb{Q} \cup \{\pm\infty\}$ -valued cost functions in general. An instance of the VCSP is a pair (\mathbb{A}, \mathbb{C}) ; its Gaifman graph, denoted by $G(\mathbb{A})$, is a graph whose vertex set is the domain A with edges between two vertices that occur together in a constraint of non-zero weight.

Minimum Solution

For minimisation, we first consider $(\mathbb{Q}_{\geq 0} \cup \{\infty\})$ -valued right-hand side structures \mathbb{C} , in which the sets of zero-valued tuples and finite-valued tuples are anti-monotone, in the following sense. There is a total order \leq_T on C , and for all tuples $\mathbf{x}, \mathbf{y} \in C^n$ with $\mathbf{x} \leq_T \mathbf{y}$ (coordinate-wise) we have that for all non-unary function symbols f of \mathbb{C} :

- $f^{\mathbb{C}}(\mathbf{x}) < \infty$ implies $f^{\mathbb{C}}(\mathbf{y}) < \infty$, and
- $f^{\mathbb{C}}(\mathbf{x}) = 0$ implies $f^{\mathbb{C}}(\mathbf{y}) = 0$.

Intuitively, larger tuples are more feasible. We call valued structures \mathbb{C} satisfying this condition Min-Sol structures. We define $\text{Min-Sol}_{\mathcal{G}}$ to be the general-valued CSP restricted to instances (\mathbb{A}, \mathbb{C}) where \mathbb{A} is a $\mathbb{Q}_{\geq 0}$ -valued structured with $G(\mathbb{A}) \in \mathcal{G}$ and \mathbb{C} is a Min-Sol structure.

For example, Weighted Minimum Vertex Cover is equivalent to the Min-Sol case where \mathbb{C} is the structure with domain $\{0, 1\}$ and $0 \leq 1$ and with a 2-ary cost function $f^{\mathbb{C}}(0, 0) = \infty$, $f^{\mathbb{C}}(1, 0) = f^{\mathbb{C}}(0, 1) = f^{\mathbb{C}}(1, 1) = 0$, and a unary cost function $u^{\mathbb{C}}(0) = 0$, $u^{\mathbb{C}}(1) = 1$.

We show that $\text{Min-Sol}_{\mathcal{G}}$ admits a PTAS for all graph classes \mathcal{G} that are efficiently Baker. (Dvořák’s definition is somewhat involved but we give an exposition in Section 4). As discussed above, this captures essentially all graph classes where a version of Baker’s technique is known to apply (including excluded-minor classes and more), except for fractionally-treewidth-fragile classes. We remark that already the very special case of Minimum Vertex Cover is not known to admit a PTAS on fractionally-treewidth-fragile classes.

Simultaneously, our results are less restrictive on the right-hand side, as unlike in earlier work such as the framework of Strict-CSP of [36], we allow arbitrary values strictly between 0 and ∞ (not only on unary constraints). Once we realise this is possible, however, the algorithm turns out to be a rather standard application of Baker’s technique: the only difference is that we increase the number of layers to account for the maximum ratio between finite, positive values (which is a constant depending on values of \mathbb{C} only).

The main novelty in our work is establishing the existence of a PTAS under a weaker assumption on the right-hand side structure \mathbb{C} – we only require that \mathbb{C} should be a diagonalisable structure. (As we will show in Section 3.5, all Min-Sol structures are diagonalisable and thus our result establishes a PTAS for Min-Sol structures as a special case.) Diagonalisability is a notion derived from the work of Brightwell and Winkler [6] in the case of graphs and Briceño, Bulatov, Dalmau, and Larose [5] in the case of relational structures (which are more general than graphs). The precise definition of diagonalisability is technical and can be found in Section 3.1. For relational structures, one characterisation is that a structure \mathbb{C} is diagonalisable if and only if the two projection homomorphisms $\pi_1, \pi_2: \mathbb{C} \times \mathbb{C} \rightarrow \mathbb{C}$ (defined as $\pi_i(x_1, x_2) = x_i$) are connected by some sequence of homomorphisms $\psi: \mathbb{C} \times \mathbb{C} \rightarrow \mathbb{C}$ such that consecutive homomorphisms in the sequence differ at only one vertex, and all the homomorphisms in the sequence are idempotent (meaning $\psi(x, x) = x$). This turns out to be equivalent to saying that for all structures \mathbb{A} , the set of all homomorphisms from

\mathbb{A} to \mathbb{C} is connected in a similar sense. A few other characterisations connect diagonalisability to statistical physics via “mixing” properties. Diagonalisability is also equivalent to finite duality (the existence of finitely many obstructions to having a homomorphism into \mathbb{C}), a notion important to the study of CSPs via logic [7]. For these and many other equivalent definitions of diagonalisability, cf. [5, Corollary 6.3 and Theorem 3.6] with $J = V(H)$.

Our main result for minimisation (proved in Section 4) is an approximation scheme for instances (\mathbb{A}, \mathbb{C}) where \mathbb{A} comes from a Baker class and \mathbb{C} is diagonalisable. One should think of the functions f_1 and f_2 as polynomials depending on the size of the graph $G(\mathbb{A})$.

THEOREM 1.1. *Let \mathcal{G} be an (f_1, f_2) -efficiently Baker class. Then, for any $\varepsilon > 0$ and any instance (\mathbb{A}, \mathbb{C}) of general-valued CSP where \mathbb{A} is a $\mathbb{Q}_{\geq 0}$ -valued structured with $G(\mathbb{A}) \in \mathcal{G}$ and \mathbb{C} is diagonalisable, we can find a solution of value at most $(1 + \varepsilon) \minval(\mathbb{A}, \mathbb{C})$ in time $f_1(|\mathbb{A}|) + f_2(c|\mathbb{A}|) \cdot c^{1/\varepsilon}$ where c depends on \mathbb{C} and \mathcal{G} only.*

Here the constant c depends polynomially on $|\mathbb{C}|$ and exponentially on the maximum ratio between certain finite positive values of \mathbb{C} . Since every class of graphs that excludes a minor is $(O(n^2), O(n))$ -efficiently Baker [16, Theorem 2.1], Section 1.1 in fact gives an EPTAS on such classes for any fixed diagonalisable structure \mathbb{C} .

Intuitively, diagonalisability allows to interpolate between any two homomorphisms, and we show this gives a natural way to combine partial solutions in the way needed in Baker’s technique (generalising the simple combination used for Vertex Cover: taking the set-theoretic sum of solutions). This proof (Theorem 3.7), which is an entirely new connection between diagonalisability and approximation, is our main contribution.

Maximum Solution

For maximisation, we extend the results of [44], which restricted the right-hand side \mathbb{C} to be $\mathbb{Q}_{\geq 0}$ -valued. We additionally allow $-\infty$ values, but the set of tuples $\mathbf{y} \in C^n$ with $f^{\mathbb{C}}(\mathbf{y}) = -\infty$ is restricted to be monotone in the following very weak sense. There is an element $c_{\perp} \in C$ such that whenever \mathbf{y} is feasible ($f^{\mathbb{C}}(\mathbf{y}) \neq -\infty$) and \mathbf{y}' is a tuple obtained from \mathbf{y} by replacing some of its elements with c_{\perp} , then \mathbf{y}' is still feasible ($f^{\mathbb{C}}(\mathbf{y}') \neq -\infty$).

We call structures \mathbb{C} satisfying this condition Max-Sol structures and we define $\text{Max-Sol}_{\mathcal{G}}$ to be the general-valued CSP restricted to instances (\mathbb{A}, \mathbb{C}) where \mathbb{A} is a $\mathbb{Q}_{\geq 0}$ -valued structured with $G(\mathbb{A}) \in \mathcal{G}$ and \mathbb{C} is a Max-Sol structure.

For example, Weighted Maximum Independent Set is equivalent to the Max-Sol case where \mathbb{C} is the structure with domain $\{0, 1\}$, with a 2-ary function $f^{\mathbb{C}}(1, 1) = -\infty$, $f^{\mathbb{C}}(0, 0) = f^{\mathbb{C}}(1, 0) = f^{\mathbb{C}}(0, 1) = 0$, and a unary function $u^{\mathbb{C}}(0) = 0$, $u^{\mathbb{C}}(1) = 1$ (so $c_{\perp} = 0$).

Our main result for maximisation (proved in Section 5) is the following.

THEOREM 1.2. *Let \mathcal{G} be a class of graphs that is fractionally-treewidth-fragile. Then $\text{Max-Sol}_{\mathcal{G}}$ admits a PTAS.*

More precisely, for all $\varepsilon > 0$, there is an algorithm that given (\mathbb{A}, \mathbb{C}) , outputs a value between $\maxval(\mathbb{A}, \mathbb{C})$ and $(1 + \varepsilon) \cdot \maxval(\mathbb{A}, \mathbb{C})$ in time $(|\mathbb{A}| + |\mathbb{C}|)^{k(\varepsilon)}$, where $k(\varepsilon)$ is a function depending on \mathcal{G} only.²

The algorithm in Theorem 1.2 does nothing more than solve a $\Theta(k(\varepsilon))$ -th level of the Sherali-Adams linear programming relaxation. This allows the algorithm to be oblivious to the graph structure, i.e. we do not assume that the fractional-treewidth-fragility of \mathcal{G} can be efficiently

²If \mathcal{G} is fractionally-tw-fragile with rate $k(\varepsilon)$ (Definition 5.16), then the exponent in the running time is $O(k(\Omega(\varepsilon)))$. See [17] for a fragility rates of various graph classes.

certified. Thus the left-hand side restriction on Gaifman graphs is the most general for which a PTAS is known; as discussed earlier, it includes excluded-minor classes and more. In fact similarly to [44], we conjecture that $\text{Max-Sol}_{\mathcal{G}}$ does not admit a PTAS for any \mathcal{G} that is not fractionally-treewidth-fragile. Since $\text{Max-Sol}_{\mathcal{G}}$ is strictly more general (by allowing negative infinite values), this conjecture might be easier to prove than the one in [44].

On the other hand, this approach does not give an EPTAS even when \mathbb{C} is fixed (i.e. the exponent of $|\mathbb{A}|$ increases with ε), and it does not construct an assignment – it only approximates the optimum value. In contrast, given a class of graphs \mathcal{G} for which fractional-treewidth-fragility can be efficiently certified (which includes essentially all known examples), it is straightforward to construct solutions to $\text{Max-Sol}_{\mathcal{G}}$ of value at least $(1 - \varepsilon) \cdot \max\text{val}(\mathbb{A}, \mathbb{C})$ in time $|\mathbb{A}| \cdot |\mathbb{C}|^{k(\varepsilon)}$.

Our main contribution in proving Theorem 1.2 is finding the right analogues of the definitions from [44] – a notion of “closeness” of structures, a dual notion that certifies this closeness with concrete mappings (a distribution of “partial homomorphisms”, see Section 5), and the proof of their equivalence (Lemma 5.7). In particular, while the name “partial homomorphism” may sound deceptively simple, we found that pin-pointing their definition (in the context of MaxSol) proved to be a surprisingly intricate balancing act.

We complement Theorem 1.2 with simple constructions which show that it is impossible to extend other results of [44] from the setting of purely optimisation Max-CSPs to the setting of general-valued CSPs , which include crisp constraints. In [44] the notion of *pliability* is defined (for Max-CSPs), which is a left-hand side restriction that takes the whole structure \mathbb{A} into account, not only its Gaifman graph, as done in this introduction so far; this allowed the authors of [44] to show that the same framework applies not only to sparse, fractionally-treewidth-fragile instances of Max-CSPs , but also to dense structures. We define an analogous notion of *strong pliability* and show in Lemma 5.13, similarly to [44], the existence of a PTAS (for general-valued CSPs) under the strong pliability assumption on the left-hand side structure, which takes the whole structure \mathbb{A} into account, not only its Gaifman graph. (Thus this is a more general tractability result than Theorem 1.2.) However, in Section C we show that even the simplest class of dense structures, namely the class of $\{0, 1\}$ -valued cliques, does not satisfy strong pliability. In fact, it is easy to show (cf. Section C) that the Max-Sol problem is hard to approximate even when the left-hand side structures are restricted to cliques.

Paper organisation. Section 2 introduces basic notations and defines the studied computational problems. The main result for minimisation, Section 1.1, is technical and proved in Section 4. In Section 3, we present the main ideas in the special case of planar structures. The main result for maximisation, Section 1.2, is proved in Section 5. Some of the proofs are deferred to Appendices A, B, and C.

2 PRELIMINARIES

For an integer k , we denote by $[k]$ the set $\{1, \dots, k\}$. For a tuple \mathbf{x} , we denote by x_i its i -th coordinate and by $\text{Set}(\mathbf{x})$ the set of elements appearing in \mathbf{x} . For two tuples \mathbf{x} and \mathbf{y} of length n , we write (\mathbf{x}, \mathbf{y}) as a shorthand for $((x_1, y_1), (x_2, y_2), \dots, (x_n, y_n))$. For a tuple \mathbf{x} of length n and a map h , we denote by $h(\mathbf{x})$ the coordinate-wise application of h ; i.e., $h(\mathbf{x}) = (h(x_1), \dots, h(x_n))$.

General-valued CSPs. A *signature* is a finite set σ of (function) symbols such as f , each with a specified arity $\text{ar}(f)$. For a set of values $\Omega \subseteq \mathbb{Q} \cup \{-\infty, +\infty\}$, an Ω -valued structure \mathbb{A} over a signature σ (or σ -structure, for short) is a finite domain A together with a function $f^{\mathbb{A}} : A^{\text{ar}(f)} \rightarrow \Omega$ for each symbol $f \in \sigma$. We denote by A, B, C, \dots the domains of structures $\mathbb{A}, \mathbb{B}, \mathbb{C}, \dots$

We define $\text{tup}(\mathbb{A})$ to be the set of all pairs (f, \mathbf{x}) such that $f \in \sigma$ and $\mathbf{x} \in A^{\text{ar}(f)}$; and $\text{tup}_{>0}(\mathbb{A})$ to be the set of all pairs $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$ with $f^{\mathbb{A}}(\mathbf{x}) > 0$.

We assume a straightforward *table encoding* of structures: the interpretation $f^{\mathbb{A}}$ of a symbol f in a structure \mathbb{A} is encoded as a collection of triples $\{(f, \mathbf{x}, f^{\mathbb{A}}(\mathbf{x})) \mid (f, \mathbf{x}) \in \text{tup}(\mathbb{A})\}$. Thus, the size of a σ -structure \mathbb{A} is roughly

$$|\mathbb{A}| = |\sigma| + |A| + \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} \log |\sigma| + \text{ar}(f) \log |A| + |\text{enc}(f^{\mathbb{A}}(\mathbf{x}))|$$

where $\text{enc}(\cdot)$ denotes a reasonable encoding for elements of \mathbb{Q} .

We consider the following computational problem.

Definition 2.1. An instance of the *general-valued CSP* (VCSP) consists of an ordered pair of σ -structures (\mathbb{A}, \mathbb{C}) . For a mapping $h: A \rightarrow C$, we define the value of h to be

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

The goal is to find the minimum or maximum value over all possible mappings $h: A \rightarrow C$, denoted $\text{minval}(\mathbb{A}, \mathbb{C})$ or $\text{maxval}(\mathbb{A}, \mathbb{C})$, respectively.

On the left-hand side we will only use $\mathbb{Q}_{\geq 0}$ -valued structures, with letters \mathbb{A}, \mathbb{B} ; on the right-hand side we will only use $\mathbb{Q}_{\geq 0} \cup \{\infty\}$ or $\mathbb{Q}_{\geq 0} \cup \{-\infty\}$ -valued structures, respectively, for minimisation and maximisation, with letters \mathbb{C}, \mathbb{D} .

For $\lambda \geq 0$ we write $\lambda\mathbb{A}$ for the *rescaled* σ -structure with domain A and $f^{\lambda\mathbb{A}}(\mathbf{x}) := \lambda f^{\mathbb{A}}(\mathbf{x})$, for $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$. For a σ -structure \mathbb{A} and subset of the domain $X \subseteq A$, we define $\mathbb{A}[X]$ to be the restriction of \mathbb{A} to X . That is, $\mathbb{A}[X]$ is a σ -structure over the domain X , and $f^{\mathbb{A}[X]}(\mathbf{x}) = f^{\mathbb{A}}(\mathbf{x})$ for each $f \in \sigma$ and $\mathbf{x} \in X^{\text{ar}(f)}$.

Following the influential work on decision CSPs by Grohe, Schwentick, and Segoufin [22], and Grohe [21], we will focus on fragments of the VCSP parametrised by the class of left-hand side structures (or their underlying class of graphs). Given a σ -structure \mathbb{A} , the *Gaifman graph* (or *primal graph*), denoted by $G(\mathbb{A})$, is the graph whose vertex set is the domain A , and whose edges are the pairs $\{u, v\}$ for which there is a tuple \mathbf{x} and a symbol $f \in \sigma$ such that u, v appear in \mathbf{x} and $f^{\mathbb{A}}(\mathbf{x}) > 0$.

For a graph parameter p and a structure \mathbb{A} , we define $p(\mathbb{A}) := p(G(\mathbb{A}))$ to be the parameter of the Gaifman graph of \mathbb{A} . In particular, the *treewidth* of \mathbb{A} is defined as $\text{tw}(\mathbb{A}) := \text{tw}(G(\mathbb{A}))$. (We will only use treewidth and excluded minors as black-boxes and thus will not need their definitions. The reader is referred to Diestel's textbook for details [13].)

Relational structures. A *relational* σ -structure \mathbb{C} includes for each symbol $f \in \sigma$ a relation $f^{\mathbb{C}} \subseteq C^{\text{ar}(f)}$. We will view relational structures as $\{0, \infty\}$ -valued structures by associating each function $f^{\mathbb{C}}: C^{\text{ar}(f)} \rightarrow \{0, \infty\}$ to the relation given by the zero-valued tuples $\{\mathbf{x} \mid f^{\mathbb{C}}(\mathbf{x}) = 0\}$. A *homomorphism* from a relational σ -structure \mathbb{C} to a relational σ -structure \mathbb{D} is a map $\psi: C \rightarrow D$ that satisfies, for every $f \in \sigma$ and every $\mathbf{x} \in C^{\text{ar}(f)}$, $f^{\mathbb{D}}(\psi(\mathbf{x})) \leq f^{\mathbb{C}}(\mathbf{x})$.

For an n -ary function f , we denote by $\text{Feas}(f)$ and $\text{Opt}(f)$ the n -ary relations defined by $\text{Feas}(f) = \{\mathbf{x} \mid f(\mathbf{x}) < \infty\}$ and $\text{Opt}(f) = \{\mathbf{x} \mid f(\mathbf{x}) = 0\}$, respectively. Let \mathbb{C} be a σ -structure. The relational σ -structure $\text{Feas}(\mathbb{C})$ contains, for each $f \in \sigma$, the relation $\text{Feas}(f^{\mathbb{C}})$; similarly, the relational σ -structure $\text{Opt}(\mathbb{C})$ contains, for every $f \in \sigma$, the relation $\text{Opt}(f^{\mathbb{C}})$.

Our results will be concerned with two particular types of right-hand side structures.

Maximum Solution

For the following definition, recall the example of Weighted Maximum Independent Set from Section 1. One should think of the element c in this context as not selecting a vertex in an independent set; the partial order on C^n then says that a subset of an independent set is also independent.

Definition 2.2 (\sqsubseteq_c). For an element c of a set C , we denote by \sqsubseteq_c the partial ordering on C defined by $c \sqsubseteq_c x$ and $x \sqsubseteq_c x$ for all $x \in C$. This induces a partial ordering on C^n coordinate-wise: we write $\mathbf{x} \sqsubseteq_c \mathbf{y}$ for $\mathbf{x}, \mathbf{y} \in C^n$ if we can obtain \mathbf{x} from \mathbf{y} by changing some (possibly none or all) of its coordinates to c .

Definition 2.3 (Max-Sol). Let σ be a finite signature. A σ -structure \mathbb{C} is called a *Max-Sol structure* if it is $(\mathbb{Q}_{\geq 0} \cup \{-\infty\})$ -valued and there is an element $c_{\perp} \in C$ such that for all $f \in \sigma$, the following holds: whenever $f^{\mathbb{C}}(\mathbf{y}) \geq 0$, we have $f^{\mathbb{C}}(\mathbf{x}) \geq 0$, for all $\mathbf{x} \sqsubseteq_{c_{\perp}} \mathbf{y}$ in $C^{\text{ar}(f)}$. Equivalently, if a tuple \mathbf{y} has non-negative value (not $-\infty$), then changing some of its coordinates to c_{\perp} still gives a non-negative value. To avoid clutter, we write \sqsubseteq_{\perp} in place of $\sqsubseteq_{c_{\perp}}$, with the choice of $c_{\perp} \in C$ implicit.

We denote by $\text{Max-Sol}_{\mathcal{G}}$ the restriction of the VCSP to instances (\mathbb{A}, \mathbb{C}) where \mathbb{A} is a $\mathbb{Q}_{\geq 0}$ -valued structure with $G(\mathbb{A}) \in \mathcal{G}$ and \mathbb{C} is a Max-Sol structure.

Observe that every $\mathbb{Q}_{\geq 0}$ -valued structure is a Max-Sol structure; thus $\text{Max-Sol}_{\mathcal{G}}$ is more general than the restriction to $\mathbb{Q}_{\geq 0}$ -valued right-hand side structures, which is the problem considered in [44].

Example 2.4. As explained in the introduction, Max-Sol structures can capture problems such as Weighted Maximum Independent Set. Another example is finding a 3-colourable induced subgraph with the maximum number of edges: take \mathbb{C} with domain $C = \{R, G, B, \perp\}$ (representing red, green, blue, as well as a fourth element corresponding to vertices not selected into the induced subgraph) and a single symbol of arity two with values $f^{\mathbb{C}}(R, R) = f^{\mathbb{C}}(G, G) = f^{\mathbb{C}}(B, B) = -\infty$, $f^{\mathbb{C}}(x, y) = 1$ for $x \neq y \in \{R, G, B\}$ and $f^{\mathbb{C}}(x, y) = 0$ otherwise. This extends to maximising the number of vertices (by introducing a unary relation), to weighted versions (by giving weights to vertices and edges of the left-hand-side structure \mathbb{A}), and to finding a maximum induced substructure satisfying an arbitrary CSP.

Remark 2.5. The “downward monotone Strict-CSP” from [36] corresponds to Definition 2.3 with some extra conditions. Firstly, there is a special unary symbol $u \in \sigma$ such that $u^{\mathbb{C}}$ is $\mathbb{Q}_{\geq 0}$ valued and all other symbols $f \in \sigma$ are $\{0, -\infty\}$ -valued (hence they express “strict” constraints). Secondly, there is a total order on C , and for each symbol $f \in \sigma$ other than u , $f^{\mathbb{C}}$ is anti-monotone; in other words, lowering some coordinates of a tuple in $C^{\text{ar}(f)}$ can not change its value from 0 to $-\infty$. (Hence the minimum element plays the role of the bottom label $c_{\perp} \in C$.)

Minimum Solution

Definition 2.6 (Min-Sol). Let σ be a finite signature. A σ -structure \mathbb{C} is called *Min-Sol* if it is $(\mathbb{Q}_{\geq 0} \cup \{\infty\})$ -valued and there is a total order \leq_{\top} on C such that: for all $f \in \sigma$ with $\text{ar}(f) > 1$ and all tuples $\mathbf{x}, \mathbf{y} \in C^n$ with $\mathbf{x} \leq_{\top} \mathbf{y}$ (coordinate-wise) we have:

- $f^{\mathbb{C}}(\mathbf{x}) < \infty$ implies $f^{\mathbb{C}}(\mathbf{y}) < \infty$, and
- $f^{\mathbb{C}}(\mathbf{x}) = 0$ implies $f^{\mathbb{C}}(\mathbf{y}) = 0$.

We denote by $\text{Min-Sol}_{\mathcal{G}}$ the restriction of the VCSP to instances (\mathbb{A}, \mathbb{C}) where \mathbb{A} is a $\mathbb{Q}_{\geq 0}$ -valued structure with $G(\mathbb{A}) \in \mathcal{G}$ and \mathbb{C} is a Min-Sol structure.

Remark 2.7. The “upward monotone Strict-CSP” from [36] corresponds to Definition 2.6 with the extra conditions that there is only one unary symbol u , $u^{\mathbb{C}}$ is monotone and injective, and all other cost functions $f^{\mathbb{C}}$ are $\{0, \infty\}$ -valued (hence they express “strict” constraints).

Remark 2.8. We observe that *some* structure (such as a total order) on the domain of a right-hand side Min-Sol structure is needed: We show how to encode 3-Colouring of planar graphs, which does not admit a PTAS (assuming $P \neq NP$).

Let G be a planar graph. Let \mathbb{A} be a structure with domain $V(G)$ over the signature $\sigma = \{u, f\}$ of arities 1 and 2, respectively. Let $u^{\mathbb{A}}(x) = 1$ for all $x \in V(G)$, and $f^{\mathbb{A}}(x, y) = 1$ if $\{x, y\} \in E(G)$ and 0 otherwise. Let \mathbb{C} be a right-hand side structure with domain $C = \{R, G, B, c_{\top}\}$. Here we think of R, G, B as three colours, and c_{\top} as a fourth extra colour we want to avoid using. We allow a monochromatic c_{\top} edge. Let $u^{\mathbb{C}}(x) = 1$ for $x = c_{\top}$ and 0 otherwise; $f^{\mathbb{C}}(R, R) = f^{\mathbb{C}}(G, G) = f^{\mathbb{C}}(B, B) = \infty$, and 0 for other pairs of values (including (c_{\top}, c_{\top})). If G is 3-colourable then $\text{minval}(\mathbb{A}, \mathbb{C}) = 0$. Otherwise, $\text{minval}(\mathbb{A}, \mathbb{C}) \geq 1$. Note that $f^{\mathbb{C}}$ respects the partial order $\sqsubseteq_{c_{\top}}$, but it does not respect any total order on C .

3 MINIMISATION ON PLANAR STRUCTURES

3.1 Diagonalisability

Briceño, Bulatov, Dalmau, and Larose defined the concepts of product structure, dismantlability, adjacency, and link graph for relational structures [5]. In this section, we will extend these concepts to valued structures in a natural way. In particular, our definitions (for structures) coincide with the definitions in [5] (for relational structures) when viewed as $\{0, \infty\}$ -valued structures.

Informally, we shall consider any two positive finite values to be basically equivalent, because we will be able to bound differences between them by a constant factor; so we shall consider an increase in value significant only if it increases from zero to positive or from finite to infinite. For a structure \mathbb{C} , we say an element $a \in C$ is *dominated* by an element $b \in C$ if we can always replace a with b : for any mapping $h: A \rightarrow C$ (from any structure \mathbb{A}), assignments to a can be changed to assignments to b without increasing $\text{val}(h)$ significantly (from zero to positive or from finite to infinite). A structure \mathbb{C} is *diagonalisable* if in the product $\mathbb{C} \times \mathbb{C}$ (defined below), one can remove dominated elements one by one until only the diagonal $\{(c, c) \mid c \in C\}$ is left. We will later see how this allows to “blend in” two different mappings h_1, h_2 from \mathbb{A} to \mathbb{C} by considering them together as a mapping to \mathbb{C}^2 .

We now proceed with formal definitions. Given two (valued) σ -structures \mathbb{C} and \mathbb{D} , we call $\psi: C \rightarrow D$ a *homomorphism* from \mathbb{C} to \mathbb{D} if ψ is a homomorphism from $\text{Feas}(\mathbb{C})$ to $\text{Feas}(\mathbb{D})$ and from $\text{Opt}(\mathbb{C})$ to $\text{Opt}(\mathbb{D})$ (in other words, finite-valued tuples map to finite-valued tuples and zero-valued tuples map to zero-valued tuples). It will be more convenient to consider both the $\text{Feas}(\mathbb{C})$ and $\text{Opt}(\mathbb{C})$ simultaneously. Thus with every structure \mathbb{C} we will associate a relational structure $\text{Rel}[\mathbb{C}]$, defined as follows.

Definition 3.1. Let σ be a valued signature. For any $f \in \sigma$, we denote by f_1 and f_2 two new relational symbols of the same arity as f . Let \mathbb{C} be a σ -structure and let $\sigma' = \bigcup_{f \in \sigma} \{f_1, f_2\}$. Define the relational σ' -structure $\text{Rel}[\mathbb{C}]$ over the domain C as follows: for each $f \in \sigma$, let $f_1^{\text{Rel}[\mathbb{C}]} = \text{Feas}(f^{\mathbb{C}}) = \{x \mid f^{\mathbb{C}}(x) < \infty\}$ and $f_2^{\text{Rel}[\mathbb{C}]} = \text{Opt}(f^{\mathbb{C}}) = \{x \mid f^{\mathbb{C}}(x) = 0\}$.

We can now define the concepts of interest for structures \mathbb{C} via the already existing concepts for relational structures $\text{Rel}[\mathbb{C}]$ from [5]. We use the following observation.

Observation 3.2. For $x, y \in \mathbb{Q}_{\geq 0} \cup \{\infty\}$, there exists $M > 0$ such that $y \leq M \cdot x$ if and only if:

- if $x < \infty$, then $y < \infty$, and
- if $x = 0$, then $y = 0$.

Given σ -structures \mathbb{C} and \mathbb{D} , we say that $\psi: C \rightarrow D$ is a *homomorphism* if ψ is a homomorphism from $\text{Rel}[\mathbb{C}]$ to $\text{Rel}[\mathbb{D}]$. Equivalently, ψ is a homomorphism if there exists $M > 0$ such that for all

$(f, \mathbf{x}) \in \text{tup}(\mathbb{C}),$

$$f^{\mathbb{D}}(\psi(\mathbf{x})) \leq M \cdot f^{\mathbb{C}}(\mathbf{x}).$$

Here we can use a uniform bound M because we only work with finite structures; it will be convenient to use this equivalent definition to keep track of the bound M .

Given σ -structures \mathbb{C} and \mathbb{D} we define the *product structure* $\mathbb{C} \times \mathbb{D}$ as a σ -structure with domain $C \times D$ and for each $f \in \sigma,$

$$f^{\mathbb{C} \times \mathbb{D}}((\mathbf{x}, \mathbf{y})) = f^{\mathbb{C}}(\mathbf{x}) + f^{\mathbb{D}}(\mathbf{y}).$$

Let $\pi_1(x, y) = x$ and $\pi_2(x, y) = y$ be the projections to the first and second coordinate, respectively. Note that π_1, π_2 are homomorphisms from \mathbb{C}^2 to \mathbb{C} for any \mathbb{C} . See Figure 1.

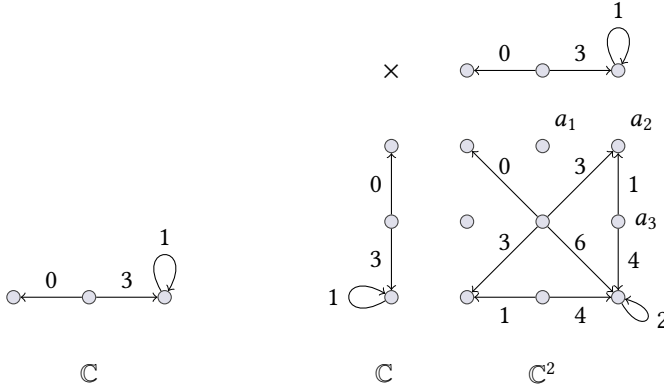


Fig. 1. Left: a digraph structure \mathbb{C} (the signature has a single symbol of arity 2) with three vertices and three arcs (tuples) with finite values – the remaining arcs have infinite values and are not drawn. Right: the product $\mathbb{C} \times \mathbb{C}$ (with three vertices a_1, a_2, a_3 distinguished for later).

We say that $a \in C$ is *dominated* by $b \in C$ if there is an $M > 0$ such that for all $(f, \mathbf{x}) \in \text{tup}(\mathbb{C})$ with $x_i = a$, we have

$$f^{\mathbb{C}}(x_1, \dots, x_{i-1}, b, x_{i+1}, \dots) \leq M \cdot f^{\mathbb{C}}(\mathbf{x}).$$

We say that $a \in C$ is *dominated in* \mathbb{C} if a is dominated by $b \neq a$ for some $b \in C$. A sequence of σ -structures $\mathbb{C}_0, \dots, \mathbb{C}_\ell$ is a *dismantling sequence* if there exists $a_i \in C_i$ such that a_i is dominated in \mathbb{C}_i , and \mathbb{C}_{i+1} is the substructure of \mathbb{C}_i induced by $C_i \setminus \{a_i\}$, for $i \in \{0, \dots, \ell - 1\}$. In this case, we say that \mathbb{C}_0 dismantles to \mathbb{C}_ℓ . A structure \mathbb{C} is *diagonalisable* if \mathbb{C}^2 dismantles to the substructure induced by its diagonal $\Delta(\mathbb{C}^2) = \{(c, c) \mid c \in C\}$.

Example 3.3. Consider \mathbb{C}^2 in Figure 1. Let f be the unique symbol (of arity two) in the signature. Let a_1, a_2, a_3 be the vertices as drawn and let b be the vertex of \mathbb{C}^2 with a loop of value 2. Then, for example, the value of the arc from a_3 to b is 4, or more formally, $f^{\mathbb{C}^2}(a_3, b) = 4$. For the vertex a_1 , all incident arcs have value ∞ (formally, $f^{\mathbb{C}^2}(a_1, x) = f^{\mathbb{C}^2}(x, a_1) = \infty$ for all $x \in V(\mathbb{C}^2)$), so it is dominated by every other vertex. The vertex a_2 is dominated by b (with $M = 4$). After removing a_2 , the vertex a_3 is dominated by b as well (this is false before removing a_2 , because $f^{\mathbb{C}^2}(a_3, a_2) = 1$ while $f^{\mathbb{C}^2}(b, a_2) = \infty$, so we cannot guarantee $f^{\mathbb{C}^2}(b, x) \leq M \cdot f^{\mathbb{C}^2}(a_3, x)$). Thus $\mathbb{C}^2, \mathbb{C}^2 \setminus \{a_1\}, \mathbb{C}^2 \setminus \{a_1, a_2\}, \mathbb{C}^2 \setminus \{a_1, a_2, a_3\}$ is a dismantling sequence. Symmetrical vertices can be similarly dominated, hence \mathbb{C}^2 dismantles to $\Delta(\mathbb{C}^2)$, meaning \mathbb{C} is diagonalisable. On the other hand, \mathbb{C} is not a MinSol structure (there is no way to order the two peripheral vertices). We refer to [6] for more examples of dismantlable graphs.

Homomorphisms ψ, ϕ from \mathbb{C} to \mathbb{D} are *adjacent* if there exists $M > 0$ such that for all $(f, \mathbf{x}) \in \text{tup}(\mathbb{C})$ and $\mathbf{y} \in D^{\text{ar}(f)}$ with $y_i \in \{\psi(x_i), \phi(x_i)\}$, we have

$$f^{\mathbb{D}}(\mathbf{y}) \leq M \cdot f^{\mathbb{C}}(\mathbf{x}). \quad (1)$$

Thus a is dominated by b in \mathbb{C} if and only if the function $s: C \rightarrow C \setminus \{a\}$ that maps a to b and everything else identically is a homomorphism from \mathbb{C} to \mathbb{C} , and s is adjacent to the identity homomorphism. (This is stronger than just s being a homomorphism, since $f^{\mathbb{C}}(a, a, a) = 0$ implies not only $f^{\mathbb{C}}(b, b, b) = 0$, but also $f^{\mathbb{C}}(a, a, b) = 0$, for example). Note that adjacency is a symmetric but not a transitive property.

Finally, for σ -structures \mathbb{C} and \mathbb{D} , we define the *link graph* $L(\mathbb{C}, \mathbb{D})$ to be the simple graph whose vertices are the homomorphisms from \mathbb{C} to \mathbb{D} , with edges between adjacent homomorphisms.

The following theorem was proved in [5, Theorem 3.6] for relational structures but the result easily extends to structures.

THEOREM 3.4. *Let \mathbb{C} be a σ -structure. Then, the following are equivalent.*

- \mathbb{C} is diagonalisable;
- π_1 and π_2 are connected in $L(\mathbb{C}^2, \mathbb{C})$ by a path of adjacent idempotent homomorphisms.
(We say a function $\psi: C^2 \rightarrow C$ is idempotent if $\psi(c, c) = c$ for all $c \in C$.)

PROOF. This follows from the fact that $\text{Rel}[\mathbb{C}^2] = \text{Rel}[\mathbb{C}]^2$ and that our definitions are the same as those of [5, Theorem 3.6] applied to $\mathbb{H} = \mathbb{C}$ and $J = \Delta(C^2)$. Specifically a function $\phi: C \rightarrow C$ is a homomorphism from \mathbb{C} to \mathbb{C} if and only if it is a homomorphism from $\text{Rel}[\mathbb{C}]$ to $\text{Rel}[\mathbb{C}]$. Similarly, a is dominated by b in \mathbb{C}^2 if and only if a is dominated by b in $\text{Rel}[\mathbb{C}^2] = \text{Rel}[\mathbb{C}]^2$. Thus \mathbb{C} is diagonalisable if and only if $\text{Rel}[\mathbb{C}]^2$ dismantles to its full diagonal (not just any subset of it). Further $\phi, \psi: C \rightarrow C$ are adjacent homomorphisms from \mathbb{C} to \mathbb{C} if and only if they are adjacent homomorphisms from $\text{Rel}[\mathbb{C}]$ to $\text{Rel}[\mathbb{C}]$. Finally, π_1, π_2 are connected by a path of adjacent idempotent homomorphisms if and only if they are J -connected by any homomorphism in $L(\text{Rel}[\mathbb{C}]^2, \mathbb{C})$ in the sense of [5]. \square

We now show that diagonalisability is more general than the Min-Sol condition.

LEMMA 3.5. *Let \mathbb{C} be a Min-Sol structure. Then \mathbb{C} is diagonalisable. Moreover, there is a path on 3 vertices between π_1 and π_2 in $L(\mathbb{C}^2, \mathbb{C})$.*

PROOF. Define $\phi: C^2 \rightarrow \Delta(C^2)$ by $\phi(x, y) = (\max(x, y), \max(x, y))$, where \max is with respect to the total order on C . We claim for each $(x, y) \in C^2$, $a := (x, y)$ is dominated by $b := \phi(x, y)$. Indeed, for each $(f, (\mathbf{x}, \mathbf{y})) \in \text{tup}(\mathbb{C}^2)$ with $(x_i, y_i) = a$ and $n := \text{ar}(f) > 1$, we have

$$\begin{aligned} & f^{\mathbb{C}^2}((x_1, y_1), \dots, b, \dots, (x_n, y_n)) \\ &= f^{\mathbb{C}}(x_1, \dots, \max(x_i, y_i), \dots, x_n) + f^{\mathbb{C}}(y_1, \dots, \max(x_i, y_i), \dots, y_n) \\ &\leq M \cdot f^{\mathbb{C}}(x_1, \dots, x_i, \dots, x_n) + M \cdot f^{\mathbb{C}}(y_1, \dots, y_i, \dots, y_n) \\ &= M \cdot f^{\mathbb{C}^2}((\mathbf{x}, \mathbf{y})), \end{aligned}$$

for some $M > 0$, where the inequality follows from the assumption that \mathbb{C} is a Min-Sol structure. For $f \in \sigma$ with $\text{ar}(f) = 1$, we have that a is dominated by b because with $M \geq 2$ we always have

$$f^{\mathbb{C}^2}(b) = f^{\mathbb{C}}(\max(x, y)) + f^{\mathbb{C}}(\max(x, y)) \leq M \cdot (f^{\mathbb{C}}(x) + f^{\mathbb{C}}(y)) = M \cdot f^{\mathbb{C}^2}(a).$$

Therefore, we can dismantle the non-diagonal elements (x, y) in any order to obtain a dismantling sequence from \mathbb{C}^2 to the substructure induced by $\Delta(C^2)$.

Let $\mu: C^2 \rightarrow C$ be defined by $\mu(x, y) = \max(x, y)$, where \max is with respect to the total order defined on C . Then similarly as above, one can check π_1, μ, π_2 is a path in $L(\mathbb{C}^2, \mathbb{C})$. \square

We remark that [5] show many other equivalent formulations, including a property known as *finite duality*. They also discuss how finite duality allows to efficiently solve many problems such as homomorphism extensions. However, in our setting it is $\text{Rel}[\mathbb{C}]$ rather than \mathbb{C} that is restricted, so such a property would not take finite, positive values of \mathbb{C} into account.

Instead, our approach is based on Baker’s technique: we partition graph into breadth-first-search layers and use the fact that the problem can be solved exactly on a subgraph induced by a few consecutive layers. To merge such solutions into one, we use a small number of overlapping layers and use the path between projections π_1, π_2 given by Theorem 3.4 to “blend in” two solutions. By increasing the number of exactly solved, non-overlapping layers, we can reduce any loss due to differences between finite, positive values.

3.2 PTAS

Baker’s approach relies on the following structural property of planar graphs, which is e.g. a direct consequence of [4, Theorem 83].

LEMMA 3.6. *Let G be a planar graph and $v_0 \in V(G)$ be an arbitrary vertex. Let L_i be the set of vertices at distance exactly i from v_0 (i.e. the i th layer of a BFS from v_0). Then, the subgraph induced by any t consecutive layers $G[L_i \cup L_{i+1} \cup \dots \cup L_{i+t+1}]$ has treewidth at most $3t$.*

THEOREM 3.7. *Let \mathcal{P} be the class of planar graphs. Then, for any $\varepsilon > 0$ and any VCSP instance (\mathbb{A}, \mathbb{C}) with $G(\mathbb{A}) \in \mathcal{P}$ and \mathbb{C} diagonalisable, we can find a solution of value at most $(1 + \varepsilon) \text{minval}(\mathbb{A}, \mathbb{C})$ in time $|\mathbb{A}| \cdot c^{1/\varepsilon}$ where c depends on \mathbb{C} only.*

PROOF. Let (\mathbb{A}, \mathbb{C}) be a VCSP instance as per the theorem. Generally, for any left-hand side structure \mathbb{B} , we will write $\text{val}_{\mathbb{B}}(h)$ for the value of an assignment $h : B \rightarrow C$ with respect to the instance (\mathbb{B}, \mathbb{C}) , and write $\text{val}(\cdot)$ for $\text{val}_{\mathbb{A}}(\cdot)$ by default.

By Theorem 3.4 there is a sequence of adjacent homomorphisms ψ_1, \dots, ψ_ℓ from \mathbb{C}^2 to \mathbb{C} such that $\psi_1 = \pi_1$ and $\psi_\ell = \pi_2$. Let $M \geq 1$ be sufficiently large such that (1) holds for all adjacent homomorphisms ψ_i and ψ_{i+1} , $i \in \{1, \dots, \ell - 1\}$. Let $k := \lceil \frac{2M}{\varepsilon} \rceil$.

Let \mathbb{A} be a $\mathbb{Q}_{\geq 0}$ -valued structure and let $G = G(\mathbb{A}) \in \mathcal{P}$ be its Gaifman graph. Fix an arbitrary vertex $v_0 \in A$ in $G(\mathbb{A})$. For $n \in \mathbb{Z}$, let $L_n \subseteq A$ be the set of vertices whose distance from v_0 is in $\{n\ell + 1, \dots, n\ell + \ell\}$. So L_n are intervals of ℓ layers, which partition the vertex set A . See Figure 2. For each $j \in \mathbb{Z}$ and $i \in [k]$ let

$$B_j^i := L_{jk-i} \cup \dots \cup L_{jk-i+k},$$




Fig. 2. Illustration of the sets L_n, B_j^i and O^i for $i = 0$ (other i look the same, just shifted).

so that B_j^i is a block of $(k+1) \cdot \ell$ consecutive layers. Iterating through the indices j gives consecutive blocks that overlap on ℓ layers; the index i shifts which layers are in the overlap. That is,

$$B_j^i \cap B_{j+1}^i = L_{(j+1)k-i}.$$

Define the overlaps $O^i = \bigcup_j B_j^i \cap B_{j+1}^i$ for $i \in [k]$. We note that the O^1, \dots, O^k are disjoint.

Consider an optimal solution $h^* : A \rightarrow C$ for the VCSP instance (\mathbb{A}, \mathbb{C}) . As the O^i are all disjoint, there exists $i^* \in [k]$ with

$$\text{val}_{\mathbb{A}[O^{i^*}]}(h^*|_{O^{i^*}}) \leq \frac{1}{k} \text{val}(h^*) \leq \frac{\varepsilon}{2M} \text{val}(h^*).$$

We henceforth write $B_j = B_j^{i^*}$ and $O = O^{i^*}$. Note that, just as in Baker's original approach, the choice of i^* is not available to the algorithm, as we do not know h^* . However, as the number of choices for $i \in [k]$ is linear in $1/\varepsilon$, we can proceed with each possible i , construct the solution h' as discussed below and output the one with the lowest value $\text{val}(h')$.

Let \mathbb{A}^+ be a σ -structure with domain A defined by

$$f^{\mathbb{A}^+}(\mathbf{x}) = \begin{cases} M \cdot f^{\mathbb{A}}(\mathbf{x}) & \text{if } \text{Set}(\mathbf{x}) \subseteq O \\ f^{\mathbb{A}}(\mathbf{x}) & \text{otherwise,} \end{cases}$$

so that tuples which lie within O are amplified by a factor of M .

For each j , the Gaifman graph $G(\mathbb{A}^+[B_j])$ has treewidth at most $O((k+1)\ell) = O(M\ell/\varepsilon)$ by (3.6). Thus for each j , we can find a tree decomposition [3] and compute an optimal solution h_j to $(\mathbb{A}^+[B_j], \mathbb{C})$ in total time $|\mathbb{A}| \cdot |\mathbb{C}|^{O(M\ell/\varepsilon)}$ via standard dynamic programming [42]. Then by optimality of h_j ,

$$\text{val}_{\mathbb{A}^+[B_j]}(h_j) \leq \text{val}_{\mathbb{A}^+[B_j]}(h^*|_{B_j}).$$

Therefore, summing over all j , we count the contribution of every constraint once, except for constraints whose scope is contained in O (and thus in exactly two sets B_j), which are counted $2M$ times in total:

$$\begin{aligned} \sum_j \text{val}_{\mathbb{A}^+[B_j]}(h_j) &\leq \sum_j \text{val}_{\mathbb{A}^+[B_j]}(h^*|_{B_j}) \\ &= \text{val}_{\mathbb{A}}(h^*) + (2M-1) \cdot \text{val}_{\mathbb{A}[O]}(h^*|_O) \leq (1+\varepsilon) \text{val}(h^*). \end{aligned} \quad (2)$$

Observe that for each $x \in A$, either $x \notin O$ and there is a unique j for which $x \in B_j$, or $x \in O$ and there is a unique j for which $x \in B_j \cap B_{j+1}$. In the latter case, $x \in L_{(j+1)k-i^*}$ and we let $s \in [\ell]$ denote the unique s for which x is at distance exactly $((j+1)k-i^*)\ell + s$ from v_0 . Let $h' : A \rightarrow C$ be defined as follows

$$h'(x) = \begin{cases} h_j(x) & \text{if } x \in B_j \text{ for a unique } j \\ \psi_s(h_j(x), h_{j+1}(x)) & \text{if } x \in B_j \cap B_{j+1} \text{ and } d(x, v_0) = ((j+1)k-i^*)\ell + s. \end{cases}$$

We claim that h' is a solution to (\mathbb{A}, \mathbb{C}) with $\text{val}(h') \leq (1+\varepsilon) \min\text{val}(\mathbb{A}, \mathbb{C})$. Let $(f, \mathbf{x}) \in \text{tup}_{>0}(\mathbb{A})$. Note that by definition of the Gaifman graph $G(\mathbb{A})$, all x_i are adjacent to each other, so $\text{Set}(\mathbf{x})$ is contained in one or two consecutive layers. Consider the following two cases.

- (a) If $\text{Set}(\mathbf{x}) \not\subseteq O$, then there is a unique j such that $\text{Set}(\mathbf{x}) \subseteq B_j$, and so $h'(x_i) = h_j(x_i)$ for each i , because either: $x_i \notin O$ and so $h'(x_i) = h_j(x_i)$ by definition of h' , or x_i is in the last layer of $B_{j-1} \cap B_j$ and $h'(x) = \psi_\ell(h_{j-1}(x), h_j(x)) = \pi_2(h_{j-1}(x), h_j(x)) = h_j(x)$, or analogously x_i is in the first layer of $B_j \cap B_{j+1}$ and $\psi_1 = \pi_1$. Thus

$$f^{\mathbb{C}}(h'(\mathbf{x})) = f^{\mathbb{C}}(h_j(\mathbf{x})). \quad (3)$$

- (b) Else, if $\text{Set}(\mathbf{x}) \subseteq O$, then there is a unique j such that $\text{Set}(\mathbf{x}) \subseteq B_j \cap B_{j+1} = L_{(j+1)k-i^*}$. Since $\text{Set}(\mathbf{x})$ is contained in two consecutive layers, there is some s such that all vertices in $\text{Set}(\mathbf{x})$ are at distance $((j+1)k - i^*)\ell + s$ or $((j+1)k - i^*)\ell + s + 1$ from v_0 . Thus

$$h(x_i) \in \{\psi_s(h_j(x_i), h_{j+1}(x_i)), \psi_{s+1}(h_j(x_i), h_{j+1}(x_i))\}$$

for each x_i . Finally, as ψ_s and ψ_{s+1} are adjacent

$$f^{\mathbb{C}}(h'(\mathbf{x})) \leq M \cdot f^{\mathbb{C}^2}(h_j(\mathbf{x}), h_{j+1}(\mathbf{x})) = M \cdot \left(f^{\mathbb{C}}(h_j(\mathbf{x})) + f^{\mathbb{C}}(h_{j+1}(\mathbf{x})) \right). \quad (4)$$

Thus, by (2), (3), and (4),

$$\text{val}_{\mathbb{A}}(h') \leq \sum_j \text{val}_{\mathbb{A}^+[B_j]}(h_j) \leq (1 + \varepsilon) \text{val}(h^*) = (1 + \varepsilon) \text{minval}(\mathbb{A}, \mathbb{C}),$$

and so h' is the solution we seek. \square

Remark 3.8. In Theorem 3.7 it would be sufficient to require that \mathbb{C}^2 dismantles to any substructure of its diagonal, as opposed to its full diagonal (as in the definition of diagonalisability). By [5, Theorem 3.6] (extended as in Theorem 3.4) this is equivalent to saying that \mathbb{C} dismantles to a substructure \mathbb{I} such that \mathbb{I} is diagonalisable.

In this case, π_1 and π_2 are still connected in $L(\mathbb{C}^2, \mathbb{C})$, but the homomorphisms in the path connecting them will not be necessarily idempotent. However, the above proof (for the case of planar graphs) did not rely on this property. This is in contrast with Theorem 1.1 (for Baker classes) where we actually use the fact that the homomorphisms are idempotent.

Since a Min-Sol structure \mathbb{C} is diagonalisable by Lemma 3.5, we have the following corollary.

COROLLARY 3.9. *Let \mathcal{P} be the class of planar graphs. Given any $\varepsilon > 0$ and instance (\mathbb{A}, \mathbb{C}) of Min-Sol $_{\mathcal{P}}$, we can find a solution of value at most $(1 + \varepsilon) \text{minval}(\mathbb{A}, \mathbb{C})$ in time $|\mathbb{A}| \cdot c^{1/\varepsilon}$, where c depends on \mathbb{C} only.*

We remark the proof yields $c^{1/\varepsilon} = |\mathbb{C}|^{O(M\ell/\varepsilon)}$, and for Min-Sol structures Lemma 3.5 yields $\ell = 3$; hence when the bound M is a constant (e.g. for $\{0, 1, \infty\}$ -valued Min-Sol structures) the dependency on \mathbb{C} is simply $|\mathbb{C}|^{O(1/\varepsilon)}$.

4 MINIMISATION ON BAKER CLASSES

4.1 Definition of Baker classes

A *layering* of a graph G is a function $\lambda: V(G) \rightarrow \mathbb{Z}$ such that $|\lambda(u) - \lambda(v)| \leq 1$ for adjacent vertices u, v in G . That is, vertices of G are partitioned into *layers* $\lambda^{-1}(i)$ for $i \in \mathbb{Z}$ and edges only go within one layer or between two consecutive layers.

Baker's technique [1] relies on a layering of planar graphs such that the subgraph induced by $\lambda^{-1}(I)$, for any interval I (a set of a few consecutive integers), has bounded treewidth (the bound depending only on $|I|$), as formally stated in Lemma 3.6. As one might imagine, this can be iterated: it would suffice that the subgraph induced by $\lambda^{-1}(I)$ itself has such a "bounded treewidth layering". Consider now the class of graphs obtained from planar graphs by adding a single vertex, adjacent to all the others; then a layering can only have three non-empty layers;³ nevertheless, an algorithm can easily circumvent this by guessing the assignment to the single new vertex (i.e. iterating through all possibilities).

Dvořák [16] defined a *Baker class* as any class of graphs that can be dealt with in the above ways. Informally, a class of graphs is Baker if any graph in the class can be reduced to an empty graph

³There can be a non-empty layer before the universal vertex, a layer containing the universal vertex, and a layer after.

by a bounded number of operations: either removing a single vertex, or selecting some layering λ and continuing separately with every interval in that layering (each subgraph induced by $\lambda^{-1}(I)$ for intervals I of at most some size). It turns out the notion of treewidth is not necessary here, as graphs of bounded treewidth also form a Baker class.

Before we state the definition formally, let us make a few remarks. The idea of iteratively going through “some layering” and then into “every interval” is conveniently phrased as a strategy winning a game in a *bounded* number of rounds (this will be particularly useful when we will want to state an assumption that the choices, including layerings, can be constructed efficiently). Dvořák’s definition considers graphs with a total ordering of their vertex set. Roughly speaking, this is to restrict the definition to “monotone” strategies, where the single vertices to be deleted are decided upfront – this restriction won’t be important for us, but we will state it as in [16] (let us also remark that layerings are not restricted by the ordering). Finally, the definition is made a bit complicated by the fact that the number of consecutive layers we may need to include in an interval may depend on how deep we go (how many iterations of the game are done). This dependency is formalised as a function r below: for a fixed problem and approximation ratio the reader should think of some arbitrarily quickly increasing function $r: \mathbb{N} \rightarrow \mathbb{N}$. We now proceed with the formal definition.

Definition 4.1. For a graph G and a function $r: \mathbb{N} \rightarrow \mathbb{N}$, the *Baker game* on (G, r) is defined as follows, for two players I and II. Player I starts by selecting a total ordering of $V(G)$. A state of the game is a pair (G', t) where G' is an induced subgraph of G (with its ordering inherited from G) and t is an integer describing how many rounds have passed. The initial state is $(G, 0)$ and Player I wins in t rounds if the state (\emptyset, t) is reached, for any t . Otherwise, in state (G', t) , Player I chooses one of the following actions:

- delete the first vertex v of G' , according to the ordering; Player II then takes no action and the game continues in state $(G' - v, t + 1)$;
- select a layering λ of G' ; Player II then selects an interval I of at most $r(t)$ (and no more than $|V(G)|$) consecutive integers and the game continues in state $(G'[\lambda^{-1}(I)], t + 1)$.

We say a function $f: \mathbb{N} \rightarrow \mathbb{N}$ is *sub-additive* if $f(x) + f(y) \leq f(x + y)$.

Definition 4.2. For sub-additive⁴ functions $f_1, f_2: \mathbb{N} \rightarrow \mathbb{N}$, we say a class of graphs \mathcal{G} is (f_1, f_2) -*efficiently Baker* if for every function $r: \mathbb{N} \rightarrow \mathbb{N}$ there exists an integer t_{\max} and an algorithm such that: for all $G \in \mathcal{G}$, the algorithm wins as Player I in the Baker game on (G, r) in at most t_{\max} rounds, using time $f_1(|G|)$ to compute the initial ordering and using time $f_2(|G|)$ to determine the action at each state of the game.

As discussed in the introduction, efficiently Baker classes generalise excluded-minor classes.

THEOREM 4.3 (DVOŘÁK [16, THEOREM 2.1]). *Let \mathcal{G} be a class of graphs that excludes a minor. Then \mathcal{G} is $(O(n^2), O(n))$ -efficiently Baker.*

4.2 PTAS

The proof for Baker classes largely follows the proof for planar graphs. The main difference is that in order to handle exceptional vertices (also known as “apex” vertices) we will need to guess and fix a partial assignment ρ on them. When “blending in” two assignments $h_j(x)$ and $h_{j+1}(x)$ to some variable x into a single assignment $\psi_s(h_j(x), h_{j+1}(x))$, we will need to preserve the partial assignment ρ , if it is defined on x . This is why we need the homomorphisms $\psi_s: \mathbb{C}^2 \rightarrow \mathbb{C}$ to be idempotent, so that $\psi_s(\rho(x), \rho(x)) = \rho(x)$.

⁴Sub-additivity is satisfied by any reasonable time-complexity bound function and is used implicitly in [16].

THEOREM (THEOREM 1.1 RESTATED). *Let \mathcal{G} be an (f_1, f_2) -efficiently Baker class. Then, for any $\varepsilon > 0$ and any instance (\mathbb{A}, \mathbb{C}) of general-valued CSP where \mathbb{A} is a $\mathbb{Q}_{\geq 0}$ -valued structured with $G(\mathbb{A}) \in \mathcal{G}$ and \mathbb{C} is diagonalisable, we can find a solution of value at most $(1 + \varepsilon) \text{minval}(\mathbb{A}, \mathbb{C})$ in time $f_1(|\mathbb{A}|) + f_2(c|\mathbb{A}|) \cdot c^{1/\varepsilon}$ where c depends on \mathbb{C} and \mathcal{G} only.*

PROOF. For a structure \mathbb{A} and an induced subgraph G' of $G(\mathbb{A})$, we write $\mathbb{A}[G']$ as a shorthand for the induced substructure $\mathbb{A}[V(G')]$ of \mathbb{A} .

Since \mathbb{C} is diagonalisable, there is a sequence ψ_1, \dots, ψ_ℓ of adjacent idempotent homomorphisms $\mathbb{C} \rightarrow \mathbb{C}$ from $\psi_1 = \pi_1$ to $\psi_\ell = \pi_2$. Let M be sufficiently large such that (1) holds for all adjacent homomorphisms ψ_i and ψ_{i+1} . For $\varepsilon > 0$, define $r(t) := 2M\ell \cdot \lceil \frac{1}{\varepsilon} \rceil \cdot 2^t$. Let Algo be the algorithm certifying that \mathcal{G} is efficiently Baker and let t_{\max} be the integer that is guaranteed to exist for r . Our algorithm starts by using Algo to compute the vertex ordering of $G(\mathbb{A})$.

We then proceed with a recursive procedure. The input of the procedure consists of an instance \mathbb{A}' , a partial assignment ρ from at most t_{\max} elements $\text{dom}(\rho)$ in \mathbb{A}' to \mathbb{C} , and a state (G', t) of the Baker game where G' is equal to the Gaifman graph of $\mathbb{A}' \setminus \text{dom}(\rho)$. We describe the procedure and claim inductively that it computes a solution of value at most $e^{\varepsilon/2^t}$ times the optimum (among all total assignments that agree with ρ on $\text{dom}(\rho)$). Moreover, we claim the procedure finishes in time at most $(2|C| \cdot M \cdot \lceil \frac{1}{\varepsilon} \rceil \cdot 2^{t_{\max}})^{t_{\max}} \cdot f_2(|G'| \cdot 2^{t_{\max}})$. Starting the recursive procedure with the instance \mathbb{A} , the state $(G(\mathbb{A}), 0)$, and empty ρ , this will conclude the proof.

Let $\text{val}(\rho)$ be the value of constraints fully contained in $\text{dom}(\rho)$:

$$\text{val}(\rho) := \sum_{\substack{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}) \\ \text{Set}(\mathbf{x}) \subseteq \text{dom}(\rho)}} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(\rho(\mathbf{x})).$$

Since all the assignment we consider will agree with ρ , $\text{val}(\rho)$ will be a common part of all of them. We will not include it in the approximation we inductively claim to get; that is, we claim the recursive procedure will output an assignment h from \mathbb{A}' to \mathbb{C} such that

$$\text{val}^\rho(h) \leq e^{\varepsilon/2^t} \cdot \text{val}^\rho(h^*),$$

where $\text{val}^\rho(h)$ is a shorthand notation for $\text{val}(h) - \text{val}(\rho)$ and h^* is an optimal solution, among solutions that agree with ρ on $\text{dom}(\rho)$.

In state (\emptyset, t) for any t , our algorithm outputs the assignment ρ , which is a total assignment since $\mathbb{A}' \setminus \text{dom}(\rho)$ is empty. It is trivially optimal (among assignments that agree with ρ).

In state (G', t) , if Algo selects to delete the minimum vertex v of G' , we consider each possible assignment to v in C , and recursively call the procedure with the state $(G' - v, t + 1)$, the instance \mathbb{A}' , and the assignment ρ extended to v . We output the best solution found this way: since we consider all possible assignments, we get a solution as close to optimum as guaranteed inductively for $t + 1$. In the final running time, the branching over all assignments to v in C will contribute at most a factor of $|C|^{t_{\max}}$.

If Algo selects a layering λ of G' , we consider each possible shift index $i \in [k]$, where $k := (2M - 1) \cdot \lceil \frac{1}{\varepsilon} \rceil \cdot 2^t$. This branching will contribute to the final running time a factor of at most $(2M \cdot \lceil \frac{1}{\varepsilon} \rceil \cdot 2^{t_{\max}})^{t_{\max}}$. For $n \in \mathbb{Z}$, let

$$L_n := \{n\ell + 1, \dots, n\ell + \ell\},$$

so L_n are intervals of length ℓ which partition \mathbb{Z} ;

$$B_j^i := L_{j-k-i} \cup \dots \cup L_{j-k+i+k},$$

so B_j^i consists of $k + 1$ such intervals, of which the last one overlaps with B_{j+1}^i :

$$O^i := \bigcup_j B_j^i \cap B_{j+1}^i = \bigcup_j L_{(j+1)k-i},$$

so O^i are disjoint for different $i \in [k]$. Note that we chose $r(t)$ to satisfy $|B_j| = (k + 1)\ell \leq r(t)$.

For a subset of integers I (such as B_j^i or O^i) we henceforth abuse notation and write $\mathbb{A}'[I]$ as a shorthand for $\mathbb{A}'[\lambda^{-1}(I) \cup \text{dom}(\rho)]$ and $h|_I$ as a shorthand for $h|_{\lambda^{-1}(I) \cup \text{dom}(\rho)}$.

Let h^* be an optimal solution to $(\mathbb{A}', \mathbb{C})$ that agrees with ρ . Since O^i are disjoint for different $i \in [k]$, so are the constraints of $\mathbb{A}'[O^i]$ (i.e. those contained in $\lambda^{-1}(O^i) \cup \text{dom}(\rho)$), except for those fully contained in $\text{dom}(\rho)$ (and accounted for in $\text{val}(\rho)$). This implies there exists a shift index $i^* \in [k]$ such that

$$\text{val}_{\mathbb{A}'[O^{i^*}]}^\rho(h^*|_{O^{i^*}}) \leq \frac{1}{k} \text{val}_{\mathbb{A}'}^\rho(h^*).$$

We henceforth only consider the recursion branch where this holds and skip the superscript i^* , and write $B_j = B_j^{i^*}$ and $O = O^{i^*}$.

Let \mathbb{A}^+ be a σ -structure with the same domain as \mathbb{A}' , defined by

$$f^{\mathbb{A}^+}(\mathbf{x}) = \begin{cases} M \cdot f^{\mathbb{A}'}(\mathbf{x}) & \text{if } \text{Set}(\mathbf{x}) \subseteq \lambda^{-1}(O) \cup \text{dom}(\rho) \\ f^{\mathbb{A}'}(\mathbf{x}) & \text{otherwise,} \end{cases}$$

so that tuples which lie within $\lambda^{-1}(O) \cup \text{dom}(\rho)$ are amplified by a factor of M .

For each j , we recurse into state $(G'[\lambda^{-1}(B_j)], t + 1)$ computing solutions h_j to $(\mathbb{A}^+[B_j], \mathbb{C})$ which agree with ρ and which by inductive assumption are almost optimal:

$$\text{val}_{\mathbb{A}^+[B_j]}^\rho(h_j) \leq e^{\varepsilon/2^{t+1}} \cdot \text{val}_{\mathbb{A}^+[B_j]}^\rho(h^*|_{B_j}).$$

Therefore, summing over j we get (by observing that every constraint is either fully contained in $\text{dom}(\rho)$, or contained in $\mathbb{A}^+[B_j]$ for exactly one j , or contained in $\mathbb{A}^+[O]$)

$$\begin{aligned} \sum_j \text{val}_{\mathbb{A}^+[B_j]}^\rho(h_j) &\leq e^{\varepsilon/2^{t+1}} \cdot \sum_j \text{val}_{\mathbb{A}^+[B_j]}^\rho(h^*|_{B_j}) \\ &= e^{\varepsilon/2^{t+1}} \cdot \left(\text{val}_{\mathbb{A}'}^\rho(h^*) + (2M - 1) \cdot \text{val}_{\mathbb{A}^+[O]}^\rho(h^*|_O) \right) \\ &\leq e^{\varepsilon/2^{t+1}} \cdot \left(1 + \frac{2M - 1}{k} \right) \text{val}_{\mathbb{A}'}^\rho(h^*) \\ &\leq e^{\varepsilon/2^t} \cdot \text{val}_{\mathbb{A}'}^\rho(h^*). \end{aligned}$$

(the last inequality holds because we chose k to satisfy $1 + \frac{2M-1}{k} \leq 1 + \frac{\varepsilon}{2^{t+1}} \leq e^{\varepsilon/2^{t+1}}$).

Observe that for each $x \in A'$ either $\lambda(x) \notin O$ there is a unique j for which $\lambda(x) \in B_j$, or $\lambda(x) \in O$ and there is a unique j for which $\lambda(x) \in B_j \cap B_{j+1} = L_{(j+1)k-i^*}$. In the latter case we let $s(x)$ denote the unique $s \in \{1, \dots, \ell\}$ for which $\lambda(x) = ((j+1)k - i^*)\ell + s$. Let $h' : A' \rightarrow C$ be defined as follows

$$h'(x) := \begin{cases} \rho(x) & \text{if } x \in \text{dom}(\rho) \\ h_j(x) & \text{if } \lambda(x) \in B_j \text{ for a unique } j \\ \psi_{s(x)}(h_j(x), h_{j+1}(x)) & \text{if } \lambda(x) \in B_j \cap B_{j+1}. \end{cases}$$

We claim that h' is a solution to $(\mathbb{A}', \mathbb{C})$ satisfying $\text{val}^\rho(h') \leq \sum_j \text{val}_{\mathbb{A}^+[B_j]}^\rho(h_j)$. This will imply

$$\text{val}_{\mathbb{A}'}^\rho(h') \leq \sum_j \text{val}_{\mathbb{A}^+[B_j]}^\rho(h_j) \leq e^{\varepsilon/2^t} \text{val}_{\mathbb{A}'}^\rho(h^*),$$

concluding that h' it is the solution we seek.

Let $(f, \mathbf{x}) \in \text{tup}_{>0}(\mathbb{A}')$. Note that since λ is a layering of $G' = G(\mathbb{A}') \setminus \text{dom}(\rho)$, there are two consecutive levels which contain all $x_i \in \text{Set}(\mathbf{x}) \setminus \text{dom}(\rho)$. Consider the three cases.

(1) If $\text{Set}(\mathbf{x}) \subseteq \text{dom}(\rho)$, then for all j ,

$$f^{\mathbb{C}}(h'(\mathbf{x})) = f^{\mathbb{C}}(\rho(\mathbf{x})) = f^{\mathbb{C}}(h_j(\mathbf{x})).$$

(2) Otherwise, if $\text{Set}(\mathbf{x}) \not\subseteq \lambda^{-1}(O) \cup \text{dom}(\rho)$, then there is a unique j such that $\text{Set}(\mathbf{x}) \subseteq \lambda^{-1}(B_j) \cup \text{dom}(\rho)$, and so $h'(x_i) = h_j(x_i)$ for each i (as some x_i might be in the first or the last layer of an overlap, but for those layers we have $\psi_1 = \pi_1$ and $\psi_\ell = \pi_2$). Thus

$$f^{\mathbb{C}}(h'(\mathbf{x})) = f^{\mathbb{C}}(h_j(\mathbf{x})).$$

(3) Else, if $\text{Set}(\mathbf{x}) \subseteq \lambda^{-1}(O) \cup \text{dom}(\rho)$ (but $\text{Set}(\mathbf{x}) \not\subseteq \text{dom}(\rho)$), then there is a unique j such that

$$\text{Set}(\mathbf{x}) \subseteq \lambda^{-1}(B_j \cap B_{j+1}) \cup \text{dom}(\rho)$$

and there is some s such that

$$\text{Set}(\mathbf{x}) \setminus \text{dom}(\rho) \subseteq \lambda^{-1}(\{(j+1)k - i^*\ell + s, ((j+1)k - i^*)\ell + s + 1\}).$$

Thus

$$h'(x_i) \in \{\psi_s(h_j(x_i), h_{j+1}(x_i)), \psi_{s+1}(h_j(x_i), h_{j+1}(x_i))\}$$

for each $x_i \in \text{Set}(\mathbf{x}) \setminus \text{dom}(\rho)$. Moreover, since ψ_s is idempotent, we can also write for $x_i \in \text{dom}(\rho)$ that

$$h'(x_i) = \rho(x_i) = \psi_s(\rho(x_i), \rho(x_i)) = \psi_s(h_j(x_i), h_{j+1}(x_i)).$$

Therefore, as ψ_t and ψ_{t+1} are adjacent, we have by definition of adjacency that

$$f^{\mathbb{C}}(h'(\mathbf{x})) \leq M \cdot f^{\mathbb{C}^2}(h_j(\mathbf{x}), h_{j+1}(\mathbf{x})) = M \cdot (f^{\mathbb{C}}(h_j(\mathbf{x})) + f^{\mathbb{C}}(h_{j+1}(\mathbf{x}))).$$

This concludes the proof that $\text{val}_{\mathbb{A}'}^{\rho}(h') \leq \sum_j \text{val}_{\mathbb{A}^+ [B_j]}^{\rho}(h_j)$ and hence h' is the solution we seek.

To check the running time, observe that every vertex of G' is contained in $\lambda^{-1}(B_j)$ for at most two j . Hence the total size of graphs $G'[\lambda^{-1}(B_j)]$ is at most $2|G'|$. Since f_2 is sub-additive, the total time required to consider those graphs in this recursive call and all sub-calls contributes a factor of at most $f_2(2^{\ell_{\max}} |G'|)$. \square

5 MAXIMISATION

To present our algorithm for maximisation, we first define what it means for two left-hand side structures \mathbb{A}, \mathbb{B} to be “close”, in a sense relevant to approximately solving Max-Sol. We then show that there is a dual view which allows to certify “closeness” by a fairly concrete mapping: a distribution of partial homomorphisms. This is then used to show that values given by Sherali-Adams linear programming relaxations of Max-Sol instances on \mathbb{A} and on \mathbb{B} are also close. Since the level- k Sherali-Adams relaxation solves the problem exactly on instances of treewidth $O(k)$, it gives a PTAS for classes of structures that are “close” to bounded treewidth, as formalised by the notion of “strong pliability” below. The proofs are similar to those in [44]; the main new contribution is finding a suitable “dual” definition (a distribution of partial homomorphisms) that makes the proofs work in the Max-Sol setting. We remark we were unable to find an analogue for the Min-Sol setting.

5.1 Pliability

Definition 5.1. For two left-hand side σ -structures \mathbb{A}, \mathbb{B} , we say \mathbb{A} *strongly overcasts* \mathbb{B} , denoted $\mathbb{A} \geq \mathbb{B}$, if for all Max-Sol σ -structures \mathbb{C} , $\text{maxval}(\mathbb{A}, \mathbb{C}) \geq \text{maxval}(\mathbb{B}, \mathbb{C})$.

In contrast, [44] defined (weak) overcasting in terms of $\mathbb{Q}_{\geq 0}$ -valued structures \mathbb{C} only, instead of the wider class of Max-Sol structures. The “strong” qualifier is only to avoid confusion with [44]: we will not consider weak overcasts in this paper, nor analogous weak variants of the definitions given below.

Definition 5.2. The *strong opt-distance* between two left-hand side σ -structures \mathbb{A} and \mathbb{B} is defined as

$$d_{\text{opt}}(\mathbb{A}, \mathbb{B}) := \inf\{\varepsilon \mid \mathbb{A} \geq e^{-\varepsilon}\mathbb{B} \text{ and } \mathbb{B} \geq e^{-\varepsilon}\mathbb{A}\}.$$

Observation 5.3. Using the fact that $\text{maxval}(\lambda\mathbb{A}, \mathbb{C}) = \lambda \text{maxval}(\mathbb{A}, \mathbb{C})$, it is an easy exercise to see that $d_{\text{opt}}(\mathbb{A}, \mathbb{B}) = \infty$ if exactly one of $\text{maxval}(\mathbb{A}, \mathbb{C})$, $\text{maxval}(\mathbb{B}, \mathbb{C})$ is $-\infty$, or exactly one of them is 0, for some Max-Sol σ -structure \mathbb{C} ; otherwise

$$d_{\text{opt}}(\mathbb{A}, \mathbb{B}) = \sup\{\varepsilon \mid \mathbb{A} \not\geq e^{-\varepsilon}\mathbb{B} \text{ or } \mathbb{B} \not\geq e^{-\varepsilon}\mathbb{A}\} = \sup_{\mathbb{C}} |\ln \text{maxval}(\mathbb{A}, \mathbb{C}) - \ln \text{maxval}(\mathbb{B}, \mathbb{C})|.$$

where the latter supremum is over all Max-Sol σ -structures \mathbb{C} such that neither is $-\infty$ nor 0. It follows that d_{opt} is symmetric and satisfies the triangle inequality.

The only graph parameter p we consider in this paper will be treewidth, tw . Just as in [44], one can prove that treedepth, or the Hadwiger number, give rise to equivalent definitions.

Definition 5.4. For a graph parameter p , a class of $\mathbb{Q}_{\geq 0}$ -valued structures \mathcal{A} is *strongly p -pliable* (with *rate* $k(\varepsilon)$) if for all $\varepsilon > 0$ there exists $k = k(\varepsilon)$ such that for every σ -structure $\mathbb{A} \in \mathcal{A}$ there exists a $\mathbb{Q}_{\geq 0}$ -valued σ -structure \mathbb{B} with $p(\mathbb{B}) \leq k$ and $d_{\text{opt}}(\mathbb{A}, \mathbb{B}) \leq \varepsilon$.

5.2 Duality

Definition 5.5 (partial functions and homomorphisms). For a partial function $g: A \rightarrow B$ and a tuple $\mathbf{x} \in A^n$, $g(\mathbf{x})$ is defined as $(g(x_1), \dots, g(x_n)) \in B^n$ if all coordinates are defined, and is undefined otherwise. For $\mathbf{y} \in B^n$, we define $g^{-1}(\mathbf{y}) := \{\mathbf{x} \in A^n \mid g(\mathbf{x}) \text{ is defined and equal to } \mathbf{y}\}$.

For left-hand side σ -structures \mathbb{A}, \mathbb{B} , a *partial homomorphism* from \mathbb{A} to \mathbb{B} is a partial function $g: A \rightarrow B$ such that: for any positive tuple $(f, \mathbf{x}) \in \text{tup}_{>0}(\mathbb{A})$, there is a positive tuple $(f, \mathbf{y}) \in \text{tup}_{>0}(\mathbb{B})$ such that $y_i = g(x_i)$ whenever $g(x_i)$ is defined (and y_i is arbitrary otherwise – in particular $y_i \neq y_j$ is allowed even if $x_i = x_j$). We denote the set of partial homomorphisms from \mathbb{A} to \mathbb{B} by $\text{p-hom}(\mathbb{A}, \mathbb{B})$.

Remark 5.6. Partial homomorphisms can also be understood as follows. For a left-hand side σ -structure \mathbb{B} , let \mathbb{B}^+ be the left-hand side σ -structure with domain $B \cup \{\star\}$, where \star is a new element, where the value for $f \in \sigma$ of arity n and an input $\mathbf{x} \in (B \cup \{\star\})^n$ is defined as

$$f^{\mathbb{B}^+}(\mathbf{x}) := \max_{\substack{\mathbf{y} \in B^n \\ \mathbf{x} \sqsubseteq_{\star} \mathbf{y}}} f^{\mathbb{B}}(\mathbf{y}).$$

In particular $f^{\mathbb{B}^+}(\mathbf{x}) = f^{\mathbb{B}}(\mathbf{x})$ for $\mathbf{x} \in B^n$. Let $\text{Pos}(\mathbb{A})$ be the relational σ -structure consisting of positive tuples of \mathbb{A} . Then a partial homomorphism g from \mathbb{A} to \mathbb{B} is the same as a homomorphism from $\text{Pos}(\mathbb{A})$ to $\text{Pos}(\mathbb{B}^+)$ (undefined assignments are the same as assignments to \star).

LEMMA 5.7. *Let \mathbb{A}, \mathbb{B} be left-hand side σ -structures. Then, the following are equivalent:*

- \mathbb{A} *strongly overcasts* \mathbb{B} , i.e. for all Max-Sol σ -structures \mathbb{C} , $\text{maxval}(\mathbb{A}, \mathbb{C}) \geq \text{maxval}(\mathbb{B}, \mathbb{C})$;

- there is a distribution of partial homomorphisms $\omega: \text{p-hom}(\mathbb{A}, \mathbb{B}) \rightarrow \mathbb{Q}_{\geq 0}$ ($\sum_g \omega(g) = 1$) such that

$$\mathbb{E}_{g \sim \omega} f^{\mathbb{A}}(g^{-1}(\mathbf{y})) \geq f^{\mathbb{B}}(\mathbf{y}) \quad \text{for all } (f, \mathbf{y}) \in \text{tup}(\mathbb{B}).$$

(Here $f^{\mathbb{A}}(g^{-1}(\mathbf{y}))$ is a shorthand for $\sum f^{\mathbb{A}}(\mathbf{x})$ over all $\mathbf{x} \in g^{-1}(\mathbf{y})$, i.e. all $\mathbf{x} \in A^{\text{ar}(f)}$ such that $g(\mathbf{x})$ is defined and equal to \mathbf{y} .)

We shall call a distribution ω from the second bullet point a *strong overcast*.

PROOF. For one direction, suppose there is a distribution ω as in the second bullet and let \mathbb{C} be a Max-Sol σ -structure with a bottom label c_{\perp} . Let $h: B \rightarrow C$ be a function achieving $\text{maxval}(\mathbb{B}, \mathbb{C})$. For $g \in \text{p-hom}(\mathbb{A}, \mathbb{B})$, let $h \circlearrowleft g: A \rightarrow C$ denote the function which maps $a \in A$ to $h(g(a))$ if $g(a)$ is defined and to c_{\perp} otherwise.

Therefore,

$$\text{maxval}(\mathbb{A}, \mathbb{C}) \geq \mathbb{E}_{g \sim \omega} \text{val}(h \circlearrowleft g) = \mathbb{E}_{g \sim \omega} \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(h \circlearrowleft g(\mathbf{x})). \quad (5)$$

We claim the expression (5) is greater or equal to

$$\text{maxval}(\mathbb{B}, \mathbb{C}) = \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{B}}(\mathbf{y}) f^{\mathbb{C}}(h(\mathbf{y})). \quad (6)$$

Indeed, suppose first that (5) is $-\infty$, or equivalently, some summand in (5) is negative. Then there exists $(f, \mathbf{x}) \in \text{tup}_{>0}(\mathbb{A})$ and $g \in \text{supp}(\omega)$ with $f^{\mathbb{C}}(h \circlearrowleft g(\mathbf{x})) = -\infty$. Since g is a partial homomorphism, by definition there exists a positive tuple $\mathbf{y} \in B^{\text{ar}(f)}$ such that $y_i = g(x_i)$ whenever $g(x_i)$ is defined. That is, $f^{\mathbb{B}}(\mathbf{y}) > 0$ and $h \circlearrowleft g(\mathbf{x}) \sqsubseteq_{\perp} h(\mathbf{y})$. By the assumption that \mathbb{C} is a Max-Sol structure, $f^{\mathbb{C}}(h(\mathbf{y})) = -\infty$. Since $f^{\mathbb{B}}(\mathbf{y})$ is positive, this gives a $-\infty$ summand in (6) and thus the inequality holds.

Otherwise, we can assume that every summand in (5) is non-negative. In that case

$$\begin{aligned} (5) &= \mathbb{E}_{g \sim \omega} \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(h \circlearrowleft g(\mathbf{x})) \\ &\geq \mathbb{E}_{g \sim \omega} \sum_{\substack{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}) \\ g(\mathbf{x}) \text{ is defined}}} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(h \circlearrowleft g(\mathbf{x})) \\ &= \mathbb{E}_{g \sim \omega} \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{A}}(g^{-1}(\mathbf{y})) f^{\mathbb{C}}(h(\mathbf{y})) \\ &= \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{C}}(h(\mathbf{y})) \mathbb{E}_{g \sim \omega} f^{\mathbb{A}}(g^{-1}(\mathbf{y})) \\ &\geq \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{C}}(h(\mathbf{y})) f^{\mathbb{B}}(\mathbf{y}) = (6), \end{aligned}$$

where after the first inequality it is still true that all summands are non-negative, and hence the last inequality follows from the assumption about ω . This concludes the proof that $\text{maxval}(\mathbb{A}, \mathbb{C}) \geq \text{maxval}(\mathbb{B}, \mathbb{C})$.

For the converse, we will use the following variant of Farkas' Lemma [43, Lemma A.2]:

Let A be an $m \times n$ rational matrix and $\bar{b} \in \mathbb{Q}^m$. Exactly one of the following holds:

- there are $x_i \in \mathbb{Q}_{\geq 0}$ ($i = 1, \dots, n$) such that $\sum_i x_i = 1$ and $\sum_i A_{i,j} x_i \geq b_j$ for $j = 1, \dots, m$;
- there are $y_j \in \mathbb{Q}_{\geq 0}$ ($j = 1, \dots, m$) such that $\sum_j A_{i,j} y_j < \sum_j b_j y_j$ for $i = 1, \dots, n$.

Suppose there is no distribution ω as in the second bullet. This means there are no numbers $\omega(g) \in \mathbb{Q}_{\geq 0}$ (for $g \in \text{p-hom}(\mathbb{A}, \mathbb{B})$) such that $\sum_g \omega(g) = 1$ and

$$\sum_{g \in \text{p-hom}(\mathbb{A}, \mathbb{B})} \omega(g) f^{\mathbb{A}}(g^{-1}(\mathbf{y})) \geq f^{\mathbb{B}}(\mathbf{y}) \quad \text{for all } (f, \mathbf{y}) \in \text{tup}(\mathbb{B}).$$

Thus by Farkas' Lemma, there are numbers $c(f, \mathbf{y}) \in \mathbb{Q}_{\geq 0}$ (for $(f, \mathbf{y}) \in \text{tup}(\mathbb{B})$) such that

$$\sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{A}}(g^{-1}(\mathbf{y})) c(f, \mathbf{y}) < \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{B}}(\mathbf{y}) c(f, \mathbf{y}) \quad \text{for all } g \in \text{p-hom}(\mathbb{A}, \mathbb{B}). \quad (7)$$

Let \mathbb{C} be the σ -structure with domain $B \cup \{c_{\perp}\}$, where c_{\perp} is a new label, and with values defined as follows for $f \in \sigma$ of arity n and $\mathbf{y} \in C^n$:

$$f^{\mathbb{C}}(\mathbf{y}) := \begin{cases} -\infty & \text{if } \forall_{\mathbf{y}' \in B^n: \mathbf{y}' \sqsubseteq_{\perp} \mathbf{y}} f^{\mathbb{B}}(\mathbf{y}') = 0 \\ & \text{(in particular if } \mathbf{y} \in B^n \text{ and } f^{\mathbb{B}}(\mathbf{y}) = 0) \\ c(f, \mathbf{y}) & \text{if } \mathbf{y} \in B^n \text{ and } f^{\mathbb{B}}(\mathbf{y}) > 0 \\ 0 & \text{otherwise; that is, if } c_{\perp} \in \mathbf{y} \text{ and } \exists_{\mathbf{y}' \in B^n: \mathbf{y}' \sqsubseteq_{\perp} \mathbf{y}} f^{\mathbb{B}}(\mathbf{y}') > 0. \end{cases}$$

We claim that \mathbb{C} is a Max-Sol structure. By Definition 2.3, we need to show that for each $f \in \sigma$, whenever $f^{\mathbb{C}}(\mathbf{y}) \geq 0$, we have $f^{\mathbb{C}}(\mathbf{x}) \geq 0$ for all $\mathbf{x} \sqsubseteq_{\perp} \mathbf{y}$. Let $f^{\mathbb{C}}(\mathbf{y}) = c(f, \mathbf{y})$ (second case in the definition of $f^{\mathbb{C}}$) and $\mathbf{x} \sqsubseteq_{\perp} \mathbf{y}$. If $\mathbf{x} \in B^n$ then $f^{\mathbb{C}}(f, \mathbf{x}) = c(f, \mathbf{x})$ and the claim holds as $c(f, \mathbf{x}) \geq 0$. If $c_{\perp} \in \mathbf{y}$ then $f^{\mathbb{C}}(\mathbf{y}) = 0$ and the claim holds again. Finally, if $f^{\mathbb{C}}(\mathbf{y}) = 0$ from the third case in the definition of $f^{\mathbb{C}}$, then for any $\mathbf{x} \sqsubseteq_{\perp} \mathbf{y}$ we have $f^{\mathbb{C}}(\mathbf{x}) = 0$ (from the third case). To show that $\text{maxval}(\mathbb{A}, \mathbb{C}) < \text{maxval}(\mathbb{B}, \mathbb{C})$, we claim that for every function g from A to $C = B \cup \{c_{\perp}\}$, we have the following strict inequality:

$$\begin{aligned} \text{val}(g) &= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(g(\mathbf{x})) = \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{C})} f^{\mathbb{A}}(g^{-1}(\mathbf{y})) f^{\mathbb{C}}(\mathbf{y}) \\ &< \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{B}}(\mathbf{y}) f^{\mathbb{C}}(\mathbf{y}) = \text{val}(\iota) \leq \text{maxval}(\mathbb{B}, \mathbb{C}), \end{aligned}$$

where ι denotes the inclusion function from B to C .

Indeed, suppose first that g , as a partial function from A to B , is not a partial homomorphism. Then there is an $(f, \mathbf{x}) \in \text{tup}_{>0}(\mathbb{A})$ such that for all $\mathbf{y} \in B^{\text{ar}(f)}$ with $\mathbf{y} \sqsubseteq_{\perp} g(\mathbf{x})$ we have $f^{\mathbb{B}}(\mathbf{y}) = 0$. Thus $f^{\mathbb{C}}(g(\mathbf{x})) = -\infty$ by definition. Thus the left-hand side of the inequality is $-\infty$, while the right-hand side is always non-negative.

Otherwise, we have

$$\begin{aligned} \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{C})} f^{\mathbb{A}}(g^{-1}(\mathbf{y})) f^{\mathbb{C}}(\mathbf{y}) &\leq \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{A}}(g^{-1}(\mathbf{y})) f^{\mathbb{C}}(\mathbf{y}) \leq \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{A}}(g^{-1}(\mathbf{y})) c(f, \mathbf{y}) \\ &< \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{B}}(\mathbf{y}) c(f, \mathbf{y}) = \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B})} f^{\mathbb{B}}(\mathbf{y}) f^{\mathbb{C}}(\mathbf{y}). \end{aligned}$$

The first inequality follows from the fact that for $(f, \mathbf{y}) \in \text{tup}(\mathbb{C}) \setminus \text{tup}(\mathbb{B})$ we have $f^{\mathbb{C}}(\mathbf{y}) \leq 0$. The second follows from $f^{\mathbb{C}}(\mathbf{y}) \leq c(f, \mathbf{y})$. The third, strict inequality follows from (7) since g (as a partial function from A to B) is a partial homomorphism. The final equality follows from the fact that if $f^{\mathbb{C}}(\mathbf{y}) \neq c(f, \mathbf{y})$ for $(f, \mathbf{y}) \in \text{tup}(\mathbb{B})$, then $f^{\mathbb{B}}(\mathbf{y}) = 0$. This concludes the proof that $\text{maxval}(\mathbb{A}, \mathbb{C}) < \text{maxval}(\mathbb{B}, \mathbb{C})$. \square

$$\begin{array}{l}
\max_{(f, \mathbf{x}) \in \text{tup } \mathbb{A}, s: \text{Set}(\mathbf{x}) \rightarrow C} \sum_{s: \text{Set}(\mathbf{x}) \rightarrow C} \lambda(\text{Set}(\mathbf{x}), s) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(s(\mathbf{x})) \\
\lambda(X, s) = \sum_{r: Y \rightarrow C, r|_X = s} \lambda(Y, r) \quad \text{for } X \subseteq Y \in \binom{A}{\leq k} \text{ and } s: X \rightarrow C \\
\sum_{s: X \rightarrow C} \lambda(X, s) = 1 \quad \text{for } X \in \binom{A}{\leq k} \\
\lambda(\text{Set}(\mathbf{x}), s) = 0 \quad \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A}) \text{ with } f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(s(\mathbf{x})) = -\infty \\
\lambda(X, s) \geq 0 \quad \text{for } X \in \binom{A}{\leq k} \text{ and } s: X \rightarrow C
\end{array}$$

Fig. 3. $\text{SA}_k(\mathbb{A}, \mathbb{C})$, the Sherali-Adams relaxation of level $k \geq \max_{f \in \sigma} \text{ar}(f)$ of (\mathbb{A}, \mathbb{C}) .

5.3 PTAS

We first define the Sherali-Adams LP hierarchy [45] for Max-Sol. Let (\mathbb{A}, \mathbb{C}) be an instance of Max-Sol over a signature σ and let $k \geq \max_{f \in \sigma} \text{ar}(f)$. We write $\binom{A}{\leq k}$ for the set of subsets of A with at most k elements. The *Sherali-Adams relaxation of level k* [45] of (\mathbb{A}, \mathbb{C}) is the linear program given in Figure 3, denoted by $\text{SA}_k(\mathbb{A}, \mathbb{C})$, which has one variable $\lambda(X, s)$ for each $X \in \binom{A}{\leq k}$ and each $s: X \rightarrow C$. We denote by $\text{maxval}_k(\mathbb{A}, \mathbb{B})$ the optimum value of $\text{SA}_k(\mathbb{A}, \mathbb{C})$, and define $\text{maxval}_k(\mathbb{A}, \mathbb{B}) = -\infty$ if $\text{SA}_k(\mathbb{A}, \mathbb{C})$ is infeasible.

Observation 5.8. Let (\mathbb{A}, \mathbb{C}) be an instance of Max-Sol, $k \geq \max_{f \in \sigma} \text{ar}(f)$ and $\lambda \geq 0$. Then, $\text{maxval}(\lambda \mathbb{A}, \mathbb{C}) = \lambda \text{maxval}(\mathbb{A}, \mathbb{C})$ and $\text{maxval}_k(\lambda \mathbb{A}, \mathbb{C}) = \lambda \text{maxval}_k(\mathbb{A}, \mathbb{C})$.

Observation 5.9. Let (\mathbb{A}, \mathbb{C}) be an instance of Max-Sol. Then, for any $k \geq \max_{f \in \sigma} \text{ar}(f)$, we have that $\text{maxval}_k(\mathbb{A}, \mathbb{C}) \geq \text{maxval}(\mathbb{A}, \mathbb{C})$.

PROOF. Let $h: A \rightarrow C$ be an optimal solution to (\mathbb{A}, \mathbb{C}) . Consider the solution $\lambda(X, s) = 1[s = h|_X]$ for $\text{SA}_k(\mathbb{A}, \mathbb{C})$. It is trivially feasible and achieves the value $\text{maxval}(\mathbb{A}, \mathbb{C})$. \square

The following easy result (proved in Section B) shows that an appropriate level of the Sherali-Adams relaxation is exact for bounded treewidth.

PROPOSITION 5.10. Let (\mathbb{A}, \mathbb{C}) be an instance of Max-Sol and $k \geq \text{tw}(\mathbb{A})$. Then, $\text{maxval}_k(\mathbb{A}, \mathbb{C}) = \text{maxval}(\mathbb{A}, \mathbb{C})$.

Definition 5.11. Let \mathbb{A} and \mathbb{B} be left-hand side σ -structures, and $k \geq \max_{f \in \sigma} \text{ar}(f)$. We write $\mathbb{A} \succeq_k \mathbb{B}$ if for all Max-Sol σ -structures \mathbb{C} we have $\text{maxval}_k(\mathbb{A}, \mathbb{C}) \geq \text{maxval}_k(\mathbb{B}, \mathbb{C})$.

Using the dual characterisation of strong overcasts (Lemma 5.7), we can show (and prove in Section A) the following.

PROPOSITION 5.12. Let \mathbb{A} and \mathbb{B} be left-hand side σ -structures, and $k \geq \max_{f \in \sigma} \text{ar}(f)$. If $\mathbb{A} \succeq \mathbb{B}$, then $\mathbb{A} \succeq_k \mathbb{B}$.

We are now ready to prove our main tractability result for maximisation problems.

LEMMA 5.13. Let \mathbb{A} be a left-hand side σ -structure, $\varepsilon \geq 0$ be small and $k \geq \max_{f \in \sigma} \text{ar}(f)$. Suppose that there exists a left-hand side σ -structure \mathbb{B} such that $\text{d}_{\text{opt}}(\mathbb{A}, \mathbb{B}) \leq \varepsilon$ and $\text{tw}(\mathbb{B}) \leq k$. Then, for every right-hand side σ -structure \mathbb{C} , we have that

$$\text{maxval}(\mathbb{A}, \mathbb{C}) \leq \text{maxval}_k(\mathbb{A}, \mathbb{C}) \leq (1 + \mathcal{O}(\varepsilon)) \text{maxval}(\mathbb{A}, \mathbb{C}).$$

PROOF. By definition of d_{opt} we have that,

$$\mathbb{A} \leq e^\varepsilon \mathbb{B} \leq e^{2\varepsilon} \mathbb{A},$$

and so

$$\mathbb{A} \leq_k e^\varepsilon \mathbb{B} \leq_k e^{2\varepsilon} \mathbb{A}$$

by Proposition 5.12. From Observations 5.8 and 5.9 and Proposition 5.10 we obtain that,

$$\text{maxval}(\mathbb{A}, \mathbb{C}) \leq \text{maxval}_k(\mathbb{A}, \mathbb{C}) \leq e^\varepsilon \text{maxval}_k(\mathbb{B}, \mathbb{C}) = e^\varepsilon \text{maxval}(\mathbb{B}, \mathbb{C}) \leq e^{2\varepsilon} \text{maxval}(\mathbb{A}, \mathbb{C}).$$

Finally, for ε small we have $e^{2\varepsilon} = 1 + O(\varepsilon)$, completing the proof. \square

Since $\text{maxval}_k(\mathbb{A}, \mathbb{C})$ can be computed in time $(|\mathbb{A}| \cdot |\mathbb{C}|)^{O(k)}$, we obtain that any strongly tw-pliable class of structures admits a PTAS.

COROLLARY 5.14. *Let \mathcal{A} be a strongly tw-pliable class of left-hand side structures. Then, the class of Max-Sol instances (\mathbb{A}, \mathbb{C}) with $\mathbb{A} \in \mathcal{A}$ admits a PTAS.*

Specifically, if \mathcal{A} is strongly tw-pliable with rate $k(\varepsilon)$, then given $\mathbb{A} \in \mathcal{A}, \mathbb{C}$, and $\varepsilon > 0$, we can output a value between $\text{maxval}(\mathbb{A}, \mathbb{C})$ and $(1 + \varepsilon) \text{maxval}(\mathbb{A}, \mathbb{C})$ in time $(|\mathbb{A}| \cdot |\mathbb{C}|)^{O(k(\Omega(\varepsilon)))}$.

In the following subsection, we show that when we look at the class of Gaifman graphs only, the appropriate condition is fractional-treewidth-fragility.

5.4 Fragility and pliability

To give Dvořák's definition of fractional fragility [15] we first define ε -thin distributions.

Definition 5.15. Let \mathcal{F} be a family of subsets of a set V and $\varepsilon > 0$. We say a distribution π over \mathcal{F} is ε -thin if $\Pr_{X \sim \pi}[v \in X] \leq \varepsilon$ for all $v \in V$.

We now give some intuition for the next definition. Consider the treewidth as a graph parameter. The idea of a *modulator*, defined below, is to remove a subset X of the vertices of a graph G to obtain a bound on the treewidth of the new graph $G - X$. The fractional variant considers a distribution over modulators. An alternative view of fractional-tw-fragility (obtained by LP duality [17]) is that for any $\varepsilon > 0$ there is k such that for any vertex weight function on G , a removal of a set vertices X whose weight is an ε -fraction of the total weight yields a graph $G - X$ of treewidth at most k .

Definition 5.16. For a graph parameter p and a number k , we define a $(p \leq k)$ -modulator of a graph G to be a set $X \subseteq V(G)$ such that $p(G - X) \leq k$. A *fractional $(p \leq k)$ -modulator* is a distribution π of such modulators X . We say that a class of graphs \mathcal{G} is *fractionally- p -fragile* (with rate $k(\varepsilon)$) if for every $\varepsilon > 0$ there is a $k = k(\varepsilon)$ such that every $G \in \mathcal{G}$ has an ε -thin fractional $(p \leq k)$ -modulator.

We need some more notation. We denote the disjoint union of graphs G and H by $G \uplus H$. For σ structures $\mathbb{A}_1, \dots, \mathbb{A}_k$, we define the σ -structure $\mathbb{B} = \uplus_{i=1}^k \mathbb{A}_i$ to be over the domain $B = \uplus_{i=1}^k A_i$ and by $f^\mathbb{B}(\mathbf{x}) = f^{\mathbb{A}_i}(\mathbf{x})$ whenever $(f, \mathbf{x}) \in \text{tup}(\mathbb{A}_i)$, and 0 otherwise.

While we are mostly interested in the following result with treewidth as the graph parameter, we state it more generally since the proof is the same.

LEMMA 5.17. *Let p be a monotone⁵ graph parameter such that $p(G \uplus H) \leq \max(p(G), p(H))$ for all graphs G and H and $p(G) \leq p(G - v) + 1$ for all $v \in V(G)$. Let \mathcal{A} be a class of structures with bounded arity r such that the class \mathcal{G} of their Gaifman graphs is fractionally- p -fragile (with rate $k(\varepsilon)$). Then \mathcal{A} is strongly p -pliable (with rate $k'(\varepsilon) = k(\Omega(\varepsilon/r)) + 1$).*

⁵ $p(H) \leq p(G)$ for all graphs G and subgraphs H of G .

The proof closely follows the proof of [43, Lemma 4.6], where the same result was shown for several particular monotone graph parameters.

PROOF. Given $\varepsilon > 0$, $\mathbb{A} \in \mathcal{A}$, let π be a fractional ($p \leq k$)-modulator such that for every $v \in V(G)$,

$$\Pr_{X \sim \pi} [v \in X] \leq \varepsilon. \quad (8)$$

For each $X \subseteq V(G) = A$ in the support of π ($\pi(X) > 0$), define $\mathbb{A}_{/X}$ to be the σ -structure obtained by contracting X to a single vertex and summing values. That is, let $\{\star_X\}$ be a new element and define $g_X : A \rightarrow (A - X) \cup \{\star_X\}$ that maps X to \star_X and $A - X$ identically. Let $\mathbb{A}_{/X}$ be over the domain $(A - X) \cup \{\star_X\}$ and

$$f^{\mathbb{A}_{/X}}(\mathbf{y}) := f^{\mathbb{A}}(g_X^{-1}(\mathbf{y})) = \sum_{\mathbf{x} \in g_X^{-1}(\mathbf{y})} f^{\mathbb{A}}(\mathbf{x})$$

for each $f \in \sigma$ of arity n and each $\mathbf{y} \in ((A - X) \cup \{\star_X\})^n$.

Define $\mathbb{B}_X = \pi(X) \cdot \mathbb{A}_{/X}$, and let $\mathbb{B} = \biguplus \mathbb{B}_X$. By definition of π and properties of p , we have $p(G(\mathbb{B}_X)) \leq p(G(\mathbb{A}) - X) + 1 \leq k + 1$, and so $p(G(\mathbb{B})) \leq k + 1$.

View g_X as a function to B (instead of as function to $B_X \subseteq B$), so that $g_X : A \rightarrow B$ is the (total) function mapping $A - X$ identically to its copy in B_X and mapping X to \star_X . It is clear that $g_X \in p\text{-hom}(\mathbb{A}, \mathbb{B})$. Define the strong overcast $\omega : \mathbb{A} \rightarrow \mathbb{B}$ to take the value g_X with probability $\pi(X)$. To check this is indeed a strong overcast, observe that for $(f, \mathbf{y}) \in \text{tup}_{>0}(\mathbb{B})$, there is a unique X such that $(f, \mathbf{y}) \in \text{tup}(\mathbb{B}_X)$, hence

$$\mathbb{E}_{g \sim \omega} f^{\mathbb{A}}(g^{-1}(\mathbf{y})) = \pi(X) f^{\mathbb{A}}(g_X^{-1}(\mathbf{y})) = f^{\mathbb{B}}(\mathbf{y}).$$

Define $g : B \rightarrow A$ to be the partial function mapping each element of $B_X - \{\star_X\}$ identically to A , leaving it undefined on \star_X . It is clear that $g \in p\text{-hom}(\mathbb{B}, \mathbb{A})$. Consider the overcast $\omega' : \mathbb{B} \rightarrow (1 - r\varepsilon)\mathbb{A}$ that is deterministically g . To check that ω' is indeed a strong overcast, let $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$. Then \mathbf{x} is covered by copies in \mathbb{B}_X for those X that do not intersect \mathbf{x} , hence

$$\begin{aligned} f^{\mathbb{B}}(g^{-1}(\mathbf{x})) &= \mathbb{E}_{X \sim \pi} [1[X \cap \mathbf{x} = \emptyset] \cdot f^{\mathbb{A}}(\mathbf{x})] \\ &= f^{\mathbb{A}}(\mathbf{x}) \Pr_{X \sim \pi} (X \cap \mathbf{x} = \emptyset) \\ &\geq f^{\mathbb{A}}(\mathbf{x}) \cdot (1 - r\varepsilon), \end{aligned}$$

where the final inequality follows by (8), the union bound, and the fact that $|\mathbf{x}| \leq r$. Hence by Lemma 5.7 applied to ω and ω' ,

$$\mathbb{A} \geq \mathbb{B} \geq (1 - r\varepsilon)\mathbb{A}.$$

By construction $p(G(\mathbb{B})) \leq k + 1$. Thus we have shown that assuming Gaifman graphs of structures in \mathcal{A} are fractionally- p -fragile with rate $k(\varepsilon)$, then for every ε and for every structure $\mathbb{A} \in \mathcal{A}$ there is a structure \mathbb{B} with $p(G(\mathbb{B})) \leq k(\varepsilon) + 1$ and $d_{\text{opt}}(\mathbb{A}, \mathbb{B}) \leq O(r\varepsilon)$. As r is fixed, this implies that \mathcal{A} is strongly p -pliable (with rate $k'(\varepsilon) = k(\Omega(\varepsilon/r)) + 1$). \square

Remark 5.18. In the above lemma, the assumption that \mathcal{A} contains structures of bounded arity r can be easily lifted, at least for $p = \text{tw}$. This is because the maximum arity of a structure in \mathcal{A} is bounded by the size of the largest clique in a Gaifman graph of a structure in \mathcal{A} . Since we assume that the class of their Gaifman graphs is fractionally- tw -fragile with rate $k(\varepsilon)$, the largest clique has at most $2k(\frac{1}{2}) + 2$ vertices (otherwise any $(\text{tw} \leq k(\frac{1}{2}))$ -modulator needs to contain more than half of the clique's vertices, and there cannot be a $\frac{1}{2}$ -thin distribution of such modulators). Thus without loss of generality we can assume $r \leq 2k(\frac{1}{2}) + 2$.

PROOF OF THEOREM 1.2. Let \mathcal{G} be a class of graphs that is fractionally-treewidth-fragile and let \mathcal{A} be a class of structures with bounded arity with Gaifman graphs in \mathcal{G} . Since treewidth satisfies the assumptions of Lemma 5.17, we have that \mathcal{A} is strongly tw-pliable. By Corollary 5.14, $\text{Max-Sol}_{\mathcal{G}}$ admits a PTAS. \square

If we only look at Gaifman graphs, one cannot use the presented approach to go beyond fractionally-treewidth-fragile classes. This is because [43, Lemma 6.1] together with the above Lemma 5.17 implies that for a class of graph \mathcal{G} and an integer r , if $\mathcal{A}_{\mathcal{G}}^{(r)}$ denotes the class of all $\mathbb{Q}_{\geq 0}$ -valued structures of arity at most r and whose Gaifman graphs are in \mathcal{G} , then $\mathcal{A}_{\mathcal{G}}^{(r)}$ is strongly tw-pliable if and only if \mathcal{G} fractionally-treewidth-fragile. In Section C, we give a simple example of a class of structures (not parametrised by their Gaifman graphs) that is strongly tw-pliable but not captured by fractional-treewidth-fragility.

ACKNOWLEDGMENTS

We would like to thank the anonymous referees of both the conference [39] and this full version of the paper. Stanislav Živný was supported by a Royal Society University Research Fellowship. Work mostly done while Balázs F. Mezei and Marcin Wrochna were employed at the University of Oxford. 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. This research was funded by UKRI EP/X024431/1.

A PROOF OF PROPOSITION 5.12

We closely follow the proof of [43, Proposition 5.3] but, given we are in a more general setting, we have to be more careful.

PROOF. Let \mathbb{C} be a Max-Sol σ -structure, and let ω be an overcast from \mathbb{A} to \mathbb{B} . Recall that for a tuple \mathbf{x} we denote by $\text{Set}(\mathbf{x})$ the set of elements appearing in \mathbf{x} . For a partial function $g: A \rightarrow B$ and a subset $X \subseteq A$, $g(X)$ denotes the set $\{g(x) \mid x \in X \text{ and } g(x) \text{ is defined}\}$. For a function $s: g(X) \rightarrow C$, $s \sqsubseteq g$ denotes the function from X to C defined as $s(g(x))$ when $g(x)$ is defined and c_{\perp} otherwise.

Let λ be an optimal solution to $\text{SA}_k(\mathbb{B}, \mathbb{C})$. That is, for each subset Y of B of size at most k , λ describes a distribution of functions to C using probabilities $\lambda(Y, s) \in \mathbb{Q}_{\geq 0}$ for $s: Y \rightarrow C$. For fixed $g \in \text{supp}(\omega)$, we define a solution λ_g to $\text{SA}_k(\mathbb{A}, \mathbb{C})$ by sampling s from this distribution and outputting $s \sqsubseteq g$. Formally, for $X \in \binom{A}{\leq k}$ and $r: X \rightarrow C$, we define

$$\lambda_g(X, r) := \sum_{s: g(X) \rightarrow C} \mathbf{1}[r = s \sqsubseteq g] \cdot \lambda(g(X), s).$$

Note that λ_g is a feasible solution. Indeed, for $X \in \binom{A}{\leq k}$ the total probability is

$$\sum_{r: X \rightarrow C} \lambda_g(X, r) = \sum_{s: g(X) \rightarrow C} \left(\sum_{r: X \rightarrow C} \mathbf{1}[r = s \sqsubseteq g] \right) \cdot \lambda(g(X), s) = \sum_{s: g(X) \rightarrow C} 1 \cdot \lambda(g(X), s) = 1;$$

while for $Z \subseteq X \in \binom{A}{\leq k}$ and $r: Z \rightarrow C$, the marginal probability of obtaining r is

$$\begin{aligned}
\sum_{\substack{t: X \rightarrow C \\ t|_Z=r}} \lambda_g(X, t) &= \sum_{\substack{t: X \rightarrow C \\ t|_Z=r}} \sum_{s: g(X) \rightarrow C} \mathbf{1}[t = s \circlearrowleft g] \cdot \lambda(g(X), s) \\
&= \sum_{s: g(X) \rightarrow C} \sum_{\substack{t: X \rightarrow C \\ t|_Z=r}} \mathbf{1}[t = s \circlearrowleft g] \cdot \lambda(g(X), s) \\
&= \sum_{s: g(X) \rightarrow C} \mathbf{1}[r = (s \circlearrowleft g)|_Z] \cdot \lambda(g(X), s) \\
&= \sum_{s: g(X) \rightarrow C} \mathbf{1}[r = (s|_{g(Z)} \circlearrowleft g)] \cdot \lambda(g(X), s) \\
&= \sum_{s': g(Z) \rightarrow C} \left(\mathbf{1}[r = s' \circlearrowleft g] \sum_{\substack{s: g(X) \rightarrow C \\ s|_{g(Z)}=s'}} \lambda(g(X), s) \right) \\
&= \sum_{s': g(Z) \rightarrow C} \mathbf{1}[r = s' \circlearrowleft g] \cdot \lambda(g(Z), s') = \lambda_g(Z, r).
\end{aligned}$$

Therefore $\max\text{val}_k(\mathbb{A}, \mathbb{C})$ is at least the expected value of the solution λ_g with g sampled from ω :

$$\begin{aligned}
\max\text{val}_k(\mathbb{A}, \mathbb{C}) &\geq \mathbb{E}_{g \sim \omega} \sum_{\substack{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}) \\ r: \text{Set}(\mathbf{x}) \rightarrow C}} \lambda_g(\text{Set}(\mathbf{x}), r) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(r(\mathbf{x})) \\
&= \mathbb{E}_{g \sim \omega} \sum_{\substack{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}) \\ s: g(\text{Set}(\mathbf{x})) \rightarrow C}} \lambda(g(\text{Set}(\mathbf{x})), s) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(s \circlearrowleft g(\mathbf{x})), \tag{9}
\end{aligned}$$

by definition of λ_g . We claim the expression (9) is at least

$$\max\text{val}_k(\mathbb{B}, \mathbb{C}) = \sum_{\substack{(f, \mathbf{y}) \in \text{tup}(\mathbb{B}) \\ s: \text{Set}(\mathbf{y}) \rightarrow C}} \lambda(\text{Set}(\mathbf{y}), s) f^{\mathbb{B}}(\mathbf{y}) f^{\mathbb{C}}(s(\mathbf{y})).$$

Indeed, suppose first that some summand in (9) is negative. Then there exist $g \in \text{supp}(\omega)$, $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$, and $s: g(\text{Set}(\mathbf{x})) \rightarrow C$ such that $\lambda(g(\text{Set}(\mathbf{x})), s) > 0$ and $f^{\mathbb{C}}(s \circlearrowleft g(\mathbf{x})) = -\infty$. Since $g \in \text{p-hom}(\mathbb{A}, \mathbb{B})$, there is some $\mathbf{y} \in A^{\text{ar}(f)}$ with $f^{\mathbb{B}}(\mathbf{y}) > 0$ such that $g(x_i)$ equals y_i whenever it is defined. In particular $g(\text{Set}(\mathbf{x})) \subseteq \text{Set}(\mathbf{y})$, hence $\lambda(\text{Set}(\mathbf{y}), s') > 0$ for some $s': \text{Set}(\mathbf{y}) \rightarrow C$ such that $s'|_{g(\text{Set}(\mathbf{x}))} = s$. This implies $s'(\mathbf{y}) \not\sqsubseteq_\perp s \circlearrowleft g(\mathbf{x})$, hence $f^{\mathbb{C}}(s'(\mathbf{y})) = -\infty$ by the assumption that \mathbb{C} is a Max-Sol structure. Therefore (A) has a summand $\lambda(\text{Set}(\mathbf{y}), s') f^{\mathbb{B}}(\mathbf{y}) f^{\mathbb{C}}(s'(\mathbf{y})) = -\infty$, so the claimed inequality (9) \geq (A) holds.

Otherwise, we can assume that every summand in (9) is non-negative. In that case

$$\begin{aligned}
(9) &= \mathbb{E}_{g \sim \omega} \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}), s: g(\text{Set}(\mathbf{x})) \rightarrow C} \lambda(g(\text{Set}(\mathbf{x})), s) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(s \circledast g(\mathbf{x})) \\
&\geq \mathbb{E}_{g \sim \omega} \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}), s: g(\text{Set}(\mathbf{x})) \rightarrow C} \lambda(g(\text{Set}(\mathbf{x})), s) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(s \circledast g(\mathbf{x})) \cdot \mathbf{1}[g(\mathbf{x}) \text{ is defined}] \\
&= \mathbb{E}_{g \sim \omega} \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B}), s: \text{Set}(\mathbf{y}) \rightarrow C} \lambda(\text{Set}(\mathbf{y}), s) f^{\mathbb{A}}(g^{-1}(\mathbf{y})) f^{\mathbb{C}}(s(\mathbf{y})) \\
&= \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B}), s: \text{Set}(\mathbf{y}) \rightarrow C} \lambda(\text{Set}(\mathbf{y}), s) \mathbb{E}_{g \sim \omega} [f^{\mathbb{A}}(g^{-1}(\mathbf{y}))] f^{\mathbb{C}}(s(\mathbf{y})) \\
&\geq \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{B}), s: \text{Set}(\mathbf{y}) \rightarrow C} \lambda(\text{Set}(\mathbf{y}), s) f^{\mathbb{B}}(\mathbf{y}) f^{\mathbb{C}}(s(\mathbf{y})) = (A),
\end{aligned}$$

where after the first inequality it is still true that all summands are non-negative, and hence the last inequality follows from the fact that ω is an overcast. This concludes the proof that $\text{maxval}_k(\mathbb{A}, \mathbb{C}) \geq \text{maxval}_k(\mathbb{B}, \mathbb{C})$. \square

B PROOF OF PROPOSITION 5.10

PROOF. We reduce to [8, Theorem 5.4], which shows that bounded treewidth implies exact solvability for minimisation of VCSPs with $\mathbb{Q}_{\geq 0} \cup \{\infty\}$ -valued right-hand side structures. We recast our relaxation into the framework of [8], which gives a more fine-grained relaxation. The SA relaxation in that paper is found in Figure 2. As we restrict to $k \geq \max_{f \in \sigma} \text{ar}(f)$, in this case variables $\lambda(f, \mathbf{x}, s)$ in [8, Figure 3] may be replaced by variables $\lambda(\text{Set}(\mathbf{x}), s)$ by equation (SA3) and the inclusion of the dummy function ρ_k in σ_k . The linear programs are now equivalent, except in [8] right-hand side σ -structures are $\mathbb{Q}_{\geq 0} \cup \{\infty\}$ -valued, and it is cast as a minimisation problem.

Let $K := \max_{(f, \mathbf{x}) \in \text{tup}(\mathbb{C})} f^{\mathbb{C}}(\mathbf{x}) \in \mathbb{Q}_{\geq 0}$. Consider a new σ -structure \mathbb{C}' defined by

$$f^{\mathbb{C}'}(\mathbf{x}) = K - f^{\mathbb{C}}(\mathbf{x}),$$

so that \mathbb{C}' is $\mathbb{Q}_{\geq 0} \cup \{\infty\}$ -valued, and thus a valued σ -structure in the framework of [8].

We have,

$$\begin{aligned}
&\max_{(f, \mathbf{x}) \in \text{tup} \mathbb{A}, s: \text{Set}(\mathbf{x}) \rightarrow C} \lambda(\text{Set}(\mathbf{x}), s) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(s(\mathbf{x})) = \\
&K \sum_{(f, \mathbf{x}) \in \text{tup} \mathbb{A}} f^{\mathbb{A}}(\mathbf{x}) - \min_{(f, \mathbf{x}) \in \text{tup} \mathbb{A}, s: \text{Set}(\mathbf{x}) \rightarrow C} \lambda(\text{Set}(\mathbf{x}), s) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}'}(s(\mathbf{x})).
\end{aligned}$$

The left term does not depend on the LP variables $\lambda(X, s)$. Thus SA_k relaxation is exact by [8, Theorem 5.4]. \square

C MAX-CSP VS. Max-Sol ON CLIQUES

In [44], the PTAS results for Max-CSPs, with non-negative rational-valued right-hand side structures, apply to many classes of dense structures as well. This is because the class of cliques (as a $\{0, 1\}$ -valued structures) and in fact any class of graphs with $\Omega(n^2)$ edges, was shown to be “tw-pliable”. In our case, because we consider more general right-hand side structures, the definition of “overcasts” and “tw-pliability” changed to “strong overcasts” and “strong tw-pliability” accordingly (simply by considering all Max-Sol right-hand side structures in place of just $\mathbb{Q}_{\geq 0}$ -valued structures). It turns out even the simplest class of dense structures, the class of cliques, is not strongly tw-pliable.

PROPOSITION C.1. *Let σ be the signature with one symbol f of arity 2. Let $\mathcal{A} = \{\mathbb{K}_n \mid n \in \mathbb{N}\}$ be the class of cliques, as σ -structures with $f^{\mathbb{K}_n}(x, y) = 1$ if $x \neq y$ and 0 otherwise, for $x, y \in [n]$. Then, \mathcal{A} is not strongly tw-pliable.*

PROOF. Suppose that \mathcal{A} is strongly tw-pliable. Let $\varepsilon = 1/10$. There exists some k such that for all n there exists a structure \mathbb{B} with $\text{tw}(\mathbb{B}) \leq k - 1$ and $d_{\text{opt}}(\mathbb{K}_n, \mathbb{B}) \leq \varepsilon$.

Let $n := 2k$ and \mathbb{B} be such that $\text{tw}(\mathbb{B}) \leq k$ and $d_{\text{opt}}(\mathbb{K}_n, \mathbb{B}) \leq \varepsilon$. It is an easy exercise that for any graph G , $\chi(G) \leq \text{tw}(G) + 1$, and therefore $\chi(G(\mathbb{B})) \leq k$, where $\chi(G)$ denotes the chromatic number of G .

Consider the following class of Max-Sol structures: for each i , let \mathbb{C}_i be a σ -structure over the domain $[i] \cup \{c_\perp\}$, and let

$$f^{\mathbb{C}_i}(x, y) = \begin{cases} -\infty & \text{if } x = y, x, y \in [i], \\ 0 & \text{if } x = y = c_\perp, \\ 1 & \text{otherwise.} \end{cases}$$

In other words, for a structure \mathbb{A} the instance $(\mathbb{A}, \mathbb{C}_i)$ asks to colour the vertices of \mathbb{A} with i colours (or assign it no colour, c_\perp), such that there are no monochromatic edges and the total weight of edges with at least one endpoint coloured is maximised. Then, $\text{maxval}(\mathbb{K}_n, \mathbb{C}_i) = i(n - 1)$ for $i \leq n$. Further, for all $i \geq \chi(G(\mathbb{B}))$, $\text{maxval}(\mathbb{B}, \mathbb{C}_i) = \text{maxval}(\mathbb{B}, \mathbb{C}_{\chi(G(\mathbb{B}))})$: the optimal solution corresponds to any proper colouring of $G(\mathbb{B})$ with colours $[i]$. As $k \geq \chi(G(\mathbb{B}))$, we have that $\text{maxval}(\mathbb{K}_n, \mathbb{C}_{2k}) / \text{maxval}(\mathbb{K}_n, \mathbb{C}_k) = 2$, but $\text{maxval}(\mathbb{B}, \mathbb{C}_{2k}) / \text{maxval}(\mathbb{B}, \mathbb{C}_k) = 1$, contradicting $d_{\text{opt}}(\mathbb{K}_n, \mathbb{B}) \leq \varepsilon$. \square

In view of Proposition C.1, one may ask whether it is possible to obtain a PTAS on cliques via different means, not relying on our notion of strong tw-pliability. It turns out that this is not possible. This follows from an easy reduction from the Maximum Clique problem.

LEMMA C.2 ([47, THEOREM 1.1]). *It is NP-hard to approximate Maximum Clique within a factor $\text{opt}^{1-\varepsilon}$ for any $\varepsilon > 0$. That is, unless $P=NP$, for any $\varepsilon > 0$ there is no polynomial-time algorithm taking a graph G and an integer r as input, that can distinguish between the following cases:*

- G has a clique of size at least r ,
- G has no clique of size r^ε .

PROPOSITION C.3. *Let $\mathcal{A} = \{\mathbb{K}_n \mid n \in \mathbb{N}\}$ be the class of cliques (as defined in Proposition C.1). Then VCSP restricted to instances (\mathbb{A}, \mathbb{C}) where $\mathbb{A} \in \mathcal{A}$ and \mathbb{C} is a Max-Sol structure, does not admit a PTAS unless $P=NP$.*

PROOF. We reduce an instance G of Maximum Clique to a suitable VCSP instance. Given a graph G on n vertices, define a σ -structure \mathbb{C} over the domain $V(G) \cup \{\star\} \cup \{c_\perp\}$ as follows. Let

$$f^{\mathbb{C}}(x, y) = \begin{cases} -\infty & \text{if } x, y \in V(G) \text{ and } xy \notin E(G), \text{ or } x = y = \star \\ 1 & \text{if } x = \star \text{ and } y \in V(G) \\ 0 & \text{otherwise.} \end{cases}$$

It is clear that \mathbb{C} is a Max-Sol structure, and it is easy to see that $\text{maxval}(\mathbb{K}_{n+1}, \mathbb{C}) = \text{Max-Clique}(G)$: in any feasible homomorphism from \mathbb{K}_{n+1} to \mathbb{C} , at most one vertex may map to each element in $V(G) \cup \{\star\}$, and vertices that are not mapped to c_\perp or \star have to map to a clique in G . The optimal solution is achieved by mapping exactly one vertex to \star , one to each vertex of a maximum clique in G and the remaining to c_\perp .

Thus by Lemma C.2, unless $P=NP$ there is no polynomial-time algorithm for approximating $\maxval(\mathbb{K}_{n+1}, \mathbb{C})$ within any sublinear factor, let alone a constant factor approximation or a PTAS. \square

Finally, we give a simple example showing that there are strongly tw-pliable (and therefore, by Lemma 5.13, tractable) classes of structures that are not captured by fractional-treewidth-fragility. Note that this does not contradict our conjecture made after Theorem 1.2 in Section 1, as the conjecture is restricted only to classes parametrised by their Gaifman graphs. The class of structures in the following proposition does *not* include all possible structures over the Gaifman graphs.

PROPOSITION C.4. *Let σ be the signature with one function symbol f of arity 2. Let \mathbb{A}_n be σ -structure over the domain $[n]$ defined by $f^{\mathbb{A}_n}(x, x) = 1$ and $f^{\mathbb{A}_n}(x, y) = 1/n$ if $x \neq y$, $x, y \in [n]$. (That is, \mathbb{A}_n is clique with loops around each vertex, loops have weight 1, and simple edges have weight $1/n$.) Then, $\mathcal{A} = \{\mathbb{A}_n \mid n \in \mathbb{N}\}$ is strongly tw-pliable but $G(\mathcal{A})$ is not fractionally-treewidth-fragile.*

PROOF. First note that $G(\mathcal{A})$ is the class of cliques with a loop on each vertex. As for any (non-empty) $X \subseteq V(K_n)$ we have $\text{tw}(K_n - X) = n - 1 - |X|$, it is easy to see that $G(\mathcal{A})$ is not fractionally-treewidth-fragile – alternatively, it follows from Proposition C.1 and Lemma 5.17 that the class of cliques is not fractionally-treewidth-fragile, and thus neither is the class of cliques with a loop around each vertex, i.e. $G(\mathcal{A})$.

We now show that \mathcal{A} is strongly tw-pliable. Let $\varepsilon > 0$ be small, $k := \lceil 2/\varepsilon \rceil$ and $n > k$ arbitrary. We show that $d_{\text{opt}}(\mathbb{A}_n, \lambda \mathbb{A}_k) \leq \varepsilon$ where $\lambda := \frac{n}{k}$.

Let ω be a random map from $V(\mathbb{A}_n) = [n]$ to $V(\mathbb{A}_k) = [k]$ and we check that it is an overcast from \mathbb{A}_n to $\lambda \mathbb{A}_k$. It is clear that ω maps positive tuples to positive tuples. Further, for $e \in E(\mathbb{A}_k)$ that is a simple edge (i.e. $e = (x, y)$ for some $x \neq y$),

$$\mathbb{E}_{g \sim \omega} f^{\mathbb{A}_n}(g^{-1}(e)) = \frac{1}{n} \binom{n}{2} \frac{2}{k^2} = \frac{n-1}{k^2} = \lambda \frac{1}{k} \cdot (1 - 1/n).$$

For $e \in E(\mathbb{A}_k)$ that is a loop (i.e. $e = (x, x)$ for some x),

$$\mathbb{E}_{g \sim \omega} f^{\mathbb{A}_n}(g^{-1}(e)) = \frac{n}{k} + \frac{1}{n} \binom{n}{2} \frac{1}{k^2} \geq \lambda.$$

As $n > \frac{2}{\varepsilon}$, therefore $\mathbb{A}_n \geq (1 - \varepsilon)\lambda \mathbb{A}_k$. By symmetry (as $k > \frac{2}{\varepsilon}$ also), $\lambda \mathbb{A}_k \geq (1 - \varepsilon)\mathbb{A}_n$. Thus, $d_{\text{opt}}(\mathbb{A}_n, \lambda \mathbb{A}_k) \leq \varepsilon$ as claimed. It follows that \mathcal{A} is strongly tw-pliable. \square

REFERENCES

- [1] Brenda S. Baker. 1994. Approximation Algorithms for NP-Complete Problems on Planar Graphs. *J. ACM* 41, 1 (1994), 153–180. <https://doi.org/10.1145/174644.174650>
- [2] Piotr Berman and Marek Karpinski. 1999. On Some Tighter Inapproximability Results (Extended Abstract). In *Proc. 26th International Colloquium on Automata, Languages and Programming (ICALP'99) (Lecture Notes in Computer Science, Vol. 1644)*. Springer, 200–209. https://doi.org/10.1007/3-540-48523-6_17
- [3] Hans L. Bodlaender. 1993. A Linear Time Algorithm for Finding Tree-Decompositions of Small Treewidth. In *Proc. 25th Annual ACM Symposium on Theory of Computing (STOC'93)*. ACM, 226–234. <https://doi.org/10.1145/167088.167161>
- [4] Hans L. Bodlaender. 1998. A Partial k -Arboretum of Graphs with Bounded Treewidth. *Theoret. Comput. Sci.* 209 (1998), 1–45. [https://doi.org/10.1016/S0304-3975\(97\)00228-4](https://doi.org/10.1016/S0304-3975(97)00228-4)
- [5] Raimundo Briceño, Andrei Bulatov, Víctor Dalmau, and Benoît Larose. 2021. Dismantlability, connectedness, and mixing in relational structures. *J. Comb. Theory, Ser. B* 147 (2021), 37–70. <https://doi.org/10.1016/j.jctb.2020.10.001> arXiv:1901.04398
- [6] Graham R. Brightwell and Peter Winkler. 2000. Gibbs Measures and Dismantlable Graphs. *J. Comb. Theory, Ser. B* 78, 1 (2000), 141–166. <https://doi.org/10.1006/jctb.1999.1935>

- [7] Andrei A. Bulatov, Andrei A. Krokhnin, and Benoît Larose. 2008. Dualities for Constraint Satisfaction Problems. In *Complexity of Constraints: An Overview of Current Research Themes (Lecture Notes in Computer Science, Vol. 5250)*, Nadia Creignou, Phokion G. Kolaitis, and Heribert Vollmer (Eds.), Springer, 93–124. https://doi.org/10.1007/978-3-540-92800-3_5
- [8] Clément Carbonnel, Miguel Romero, and Stanislav Živný. 2022. The Complexity of General-Valued Constraint Satisfaction Problems Seen from the Other Side. *SIAM J. Comput.* 51, 1 (2022), 19–69. <https://doi.org/10.1137/19m1250121> arXiv:1710.03148
- [9] Martin C. Cooper and Stanislav Živný. 2017. Hybrid Tractable Classes of Constraint Problems. In *The Constraint Satisfaction Problem: Complexity and Approximability*, Andrei A. Krokhnin and Stanislav Živný (Eds.), Dagstuhl Follow-Ups, Vol. 7. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 113–135. <https://doi.org/10.4230/DFU.Vol7.15301.4>
- [10] David P. Dailey. 1980. Uniqueness of colorability and colorability of planar 4-regular graphs are NP-complete. *Discrete Math.* 30, 3 (1980), 289–293. [https://doi.org/10.1016/0012-365X\(80\)90236-8](https://doi.org/10.1016/0012-365X(80)90236-8)
- [11] Erik D. Demaine, Mohammad Taghi Hajiaghayi, and Ken-ichi Kawarabayashi. 2005. Algorithmic Graph Minor Theory: Decomposition, Approximation, and Coloring. In *Proc. 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. IEEE Computer Society, 637–646. <https://doi.org/10.1109/SFCS.2005.14>
- [12] Matt DeVos, Guoli Ding, Bogdan Oporowski, Daniel P. Sanders, Bruce Reed, Paul Seymour, and Dirk Vertigan. 2004. Excluding any graph as a minor allows a low tree-width 2-coloring. *J. Comb. Theory B* 91, 1 (2004), 25–41. <https://doi.org/10.1016/j.jctb.2003.09.001>
- [13] Reinhard Diestel. 2010. *Graph Theory* (fourth ed.). Springer.
- [14] Zdeněk Dvořák. [n.d.]. Personal communication. One construction is as follows: start with an arbitrarily large integer m and for i from m down to 1, introduce an independent set of i new vertices and connect them via paths of length i to all previous vertices.
- [15] Zdeněk Dvořák. 2016. Sublinear separators, fragility and subexponential expansion. *European J. Combin.* 52 (2016), 103–119. <https://doi.org/10.1016/j.ejc.2015.09.001> arXiv:1404.7219
- [16] Zdeněk Dvořák. 2020. Baker game and polynomial-time approximation schemes. In *Proc. 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'20)*. SIAM, 2227–2240. <https://doi.org/10.1137/1.9781611975994.137> arXiv:1901.01797
- [17] Zdenek Dvořák and Jean-Sébastien Sereni. 2020. On Fractional Fragility Rates of Graph Classes. *Electron. J. Comb.* 27, 4 (2020), P4.9. arXiv:1907.12634 <https://www.combinatorics.org/ojs/index.php/eljc/article/view/v27i4p9>
- [18] Peter Fulla, Hannes Uppman, and Stanislav Živný. 2019. The Complexity of Boolean Surjective General-Valued CSPs. *ACM Trans. Comput. Theory* 11, 1 (2019), 4:1–4:31. <https://doi.org/10.1145/3282429> arXiv:1702.04679
- [19] Alexander Grigoriev and Hans L. Bodlaender. 2007. Algorithms for Graphs Embeddable with Few Crossings per Edge. *Algorithmica* 49, 1 (01 9 2007), 1–11. <https://doi.org/10.1007/s00453-007-0010-x>
- [20] Martin Grohe. 2003. Local Tree-Width, Excluded Minors, and Approximation Algorithms. *Combinatorica* 23, 4 (2003), 613–632. <https://doi.org/10.1007/s00493-003-0037-9>
- [21] Martin Grohe. 2007. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM* 54, 1 (2007), 1–24. <https://doi.org/10.1145/1206035.1206036>
- [22] Martin Grohe, Thomas Schwentick, and Luc Segoufin. 2001. When is the evaluation of conjunctive queries tractable?. In *Proc. 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*. ACM, 657–666. <https://doi.org/10.1145/380752.380867>
- [23] Gregory Gutin, Pavol Hell, Arash Rafiey, and Anders Yeo. 2008. A dichotomy for minimum cost graph homomorphisms. *European J. Combin.* 29, 4 (2008), 900–911. <https://doi.org/10.1016/j.ejc.2007.11.012>
- [24] Pavol Hell, Monaldo Mastrolilli, Mayssam Mohammadi Nevisi, and Arash Rafiey. 2012. Approximation of Minimum Cost Homomorphisms. In *Proc. 20th Annual European Symposium on Algorithms (ESA12) (Lecture Notes in Computer Science, Vol. 7501)*. Springer, 587–598. https://doi.org/10.1007/978-3-642-33090-2_51
- [25] Pavol Hell and Arash Rafiey. 2012. The Dichotomy of Minimum Cost Homomorphism Problems for Digraphs. *SIAM J. Discrete Math* 26, 4 (2012), 1597–1608. <https://doi.org/10.1137/100783856>
- [26] Harry B. Hunt III, Madhav V. Marathe, Venkatesh Radhakrishnan, S.S Ravi, Daniel J. Rosenkrantz, and Richard E. Stearns. 1998. NC-Approximation Schemes for NP- and PSPACE-Hard Problems for Geometric Graphs. *J. Algorithms* 26, 2 (1998), 238–274. <https://doi.org/10.1006/jagm.1997.0903>
- [27] Peter Jeavons, Andrei A. Krokhnin, and Stanislav Živný. 2014. The Complexity of Valued Constraint Satisfaction. *Bull. EATCS* 113 (2014). <http://eatcs.org/beatcs/index.php/beatcs/article/view/266>
- [28] Peter Jonsson, Andrei A. Krokhnin, and Fredrik Kuivinen. 2009. Hard constraint satisfaction problems have hard gaps at location 1. *Theor. Comput. Sci.* 410, 38–40 (2009), 3856–3874. <https://doi.org/10.1016/j.tcs.2009.05.022>
- [29] Peter Jonsson, Fredrik Kuivinen, and Gustav Nordh. 2008. MAX ONES Generalized to Larger Domains. *SIAM J. Comput.* 38, 1 (2008), 329–365. <https://doi.org/10.1137/060669231>
- [30] Peter Jonsson and Gustav Nordh. 2008. Introduction to the Maximum Solution Problem. In *Complexity of Constraints: An Overview of Current Research Themes (Lecture Notes in Computer Science, Vol. 5250)*, Nadia Creignou, Phokion G.

- Kolaitis, and Heribert Vollmer (Eds.). Springer, 255–282. https://doi.org/10.1007/978-3-540-92800-3_10
- [31] Sanjeev Khanna and Rajeev Motwani. 1996. Towards a Syntactic Characterization of PTAS. In *Proc. 28th Annual ACM Symposium on the Theory of Computing (STOC'96)*. ACM, 329–337. <https://doi.org/10.1145/237814.237979>
- [32] Sanjeev Khanna, Madhu Sudan, Luca Trevisan, and David P. Williamson. 2001. The Approximability of Constraint Satisfaction Problems. *SIAM J. Comput.* 30, 6 (2001), 1863–1920. <https://doi.org/10.1137/S0097539799349948>
- [33] Subhash Khot. 2002. On the power of unique 2-prover 1-round games. In *Proc. 34th Annual ACM Symposium on Theory of Computing (STOC'02)*. ACM, 767–775. <https://doi.org/10.1145/509907.510017>
- [34] Vladimir Kolmogorov, Andrei A. Krokhin, and Michal Rolinek. 2017. The Complexity of General-Valued CSPs. *SIAM J. Comput.* 46, 3 (2017), 1087–1110. <https://doi.org/10.1137/16M1091836> arXiv:1502.07327
- [35] Marcin Kozik and Joanna Ochremiak. 2015. Algebraic Properties of Valued Constraint Satisfaction Problem. In *Proc. 42nd International Colloquium on Automata, Languages, and Programming (ICALP'15) (Lecture Notes in Computer Science, Vol. 9134)*. Springer, 846–858. https://doi.org/10.1007/978-3-662-47672-7_69
- [36] Amit Kumar, Rajsekar Manokaran, Madhur Tulsiani, and Nisheeth K Vishnoi. 2011. On LP-based approximability for strict CSPs. In *Proc. 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*. SIAM, 1560–1573. <https://doi.org/10.1137/1.9781611973082.121>
- [37] Konstantin Makarychev and Yury Makarychev. 2017. Approximation Algorithms for CSPs. In *The Constraint Satisfaction Problem: Complexity and Approximability*, Andrei A. Krokhin and Stanislav Živný (Eds.). Dagstuhl Follow-Ups, Vol. 7. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 287–325. <https://doi.org/10.4230/DFU.Vol7.15301.11>
- [38] Madhav V. Marathe, Harry B. Hunt III, and Richard E. Stearns. 1997. Level-treewidth property, exact algorithms and approximation schemes. In *Proc. 29th Annual ACM Symposium on Theory of Computing (STOC'97)*. ACM. <https://www.osti.gov/biblio/471394>
- [39] Balázs F. Mezei, Marcin Wrochna, and Stanislav Živný. 2021. PTAS for sparse general-valued CSPs. In *Proc. 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'21)*. IEEE, 1–11. <https://doi.org/10.1109/LICS52264.2021.9470599> arXiv:2012.12607
- [40] Akbar Rafiey, Arash Rafiey, and Thiago Santos. 2019. Toward a Dichotomy for Approximation of H-Coloring. In *Proc. 46th International Colloquium on Automata, Languages, and Programming (ICALP'19) (LIPIcs, Vol. 132)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 91:1–91:16. <https://doi.org/10.4230/LIPIcs.ICALP.2019.91> arXiv:1902.02201
- [41] Prasad Raghavendra. 2008. Optimal algorithms and inapproximability results for every CSP?. In *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC'08)*. ACM, 245–254. <https://doi.org/10.1145/1374376.1374414>
- [42] Bruce A Reed. 2003. Algorithmic aspects of tree width. In *Recent advances in algorithms and combinatorics*. Springer, 85–107.
- [43] Miguel Romero, Marcin Wrochna, and Stanislav Živný. 2020. *Treewidth-Pliability and PTAS for Max-CSPs*. Technical Report. arXiv:1911.03204
- [44] Miguel Romero, Marcin Wrochna, and Stanislav Živný. 2021. Treewidth-Pliability and PTAS for Max-CSPs. In *Proc. 2021 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'21)*. SIAM, 473–483. <https://doi.org/10.1137/1.9781611976465.29> arXiv:1911.03204
- [45] Hanif D. Sherali and Warren P. Adams. 1990. A Hierarchy of Relaxations between the Continuous and Convex Hull Representations for Zero-One Programming Problems. *SIAM J. Discrete Math.* 3, 3 (1990), 411–430. <https://doi.org/10.1137/0403036>
- [46] Rustem Takhanov. 2010. A Dichotomy Theorem for the General Minimum Cost Homomorphism Problem. In *Proc. 27th International Symposium on Theoretical Aspects of Computer Science (STACS'10) (LIPIcs)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 657–668. <https://doi.org/10.4230/LIPIcs.STACS.2010.2493>
- [47] David Zuckerman. 2007. Linear Degree Extractors and the Inapproximability of Max Clique and Chromatic Number. *Theory Comput.* 3, 1 (2007), 103–128. <https://doi.org/10.4086/toc.2007.v003a006>