PTAS for Sparse General-Valued CSPs

Balázs F. Mezei
University of Oxford, UK
balazs.mezei@cs.ox.ac.uk

Marcin Wrochna
University of Oxford, UK
m.wrochna@gmail.com

Stanislav Živný
University of Oxford, UK
standa.zivny@cs.ox.ac.uk

***Abstract*—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.

I. Introduction

Min-Ones and Max-Ones, studied by Khanna and Motwani (under the names of TMIN and TMAX, respectively) [1] and by Khanna, Sudan, Trevisan, and Williamson [2], 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 Kumar, Manokaran, Tulsiani, and Vishnoi [3], and Min/Max-Solution by Jonsson, Kuivinen, and Nordh [4]. With arbitrary unary functions, such problems have been studied under the name of Min-Cost-Hom by Gutin, Hell, Rafiey, and Yeo [5], Takhanov [6], and others [7], [8], [9]. In this paper we consider the still more general setting of *general-valued* 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 [10], [11], [12] for surveys), we seek the most general conditions that allow to obtain a polynomial-time approximation scheme (PTAS).

Baker [13] 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 [14], [15], [16] and beyond: e.g. graphs embeddable in a fixed surface with few intersections per edge [17], [18], or sparse unit ball intersection graphs in few dimensions [19] (but not e.g. 3-regular expanders: bounded degree is not sufficient to get a PTAS even for Independent Set [20]). Dvořák [21] 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 [22] 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

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

978-1-6654-4895-6/21/$31.00

the examples discussed above; on the other hand, it is strictly included in the family of fractionally-treewidth-fragile classes (and hence less general) [23]. 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 [22] 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 [24]. 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 [25]. 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" 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 [29] 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.

A. Related work

The exact solvability of general-valued CSPs has been characterised for left-hand side restrictions [30] (tractable cases are precisely classes that have bounded treewidth, up to a certain notion of homomorphic equivalence) and right-hand side restrictions [31] (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. [2], 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 [4] and [10] 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 [32]. 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) [34]. Analogous results for monotone Strict-CSPs were obtained by Kumar et al. [3]. Constant-factor approximation algorithms have been established for right-hand side restricted Min-Cost-Hom on special graphs by Hell, Mastrolilli, Nevisi, and Rafiey [8], and for all graphs and some digraphs by Rafiey, Rafiey, and Santos [9].

B. Our results

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 II. Briefly, a *valued structure* $\mathbb{A}$ consists of a domain A and a collection of functions $f^{\mathbb{A}}\colon A^n \to \mathbb{Q} \cup \{\pm\infty\}$, indexed by symbols f belonging to a set of symbols σ called a signature. For two structures $\mathbb{A}, \mathbb{C}$, the *value* of an assignment $h\colon A \to C$ 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 $\mathrm{minval}(\mathbb{A}, \mathbb{C})$ and $\mathrm{maxval}(\mathbb{A}, \mathbb{C})$ 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 $\mathrm{G}(\mathbb{A})$, 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 $\leq_{\top}$ on C, and for

[1]This follows from the NP-hardness result of Kozik and Ochremiak [26], which actually shows APX-hardness; for earlier, explicit APX-hardness results for CSPs see, e.g., [4], [27]. However, we remark non-trivial PTAS examples are known for "surjective" maximisation finite-valued CSPs [28].

all tuples $\mathbf{x}, \mathbf{y} \in C^n$ with $\mathbf{x} \leq_\top \mathbf{y}$ (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 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 $\mathrm{G}(\mathbb{A}) \in \mathcal{G}$ 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 $\{0,1\}$ 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 $u^{\mathbb{C}}(0) = 0$, $u^{\mathbb{C}}(1) = 1$.

In the full version of this paper [35], we show that Min-Sol$_{\mathcal{G}}$ admits a PTAS for all graph classes $\mathcal{G}$ that are efficiently Baker. 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 [3], 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 Lemma III.4, 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 [36] in the case of graphs and Briceño, Bulatov, Dalmua, and Larose [37] 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 III-A. For relational structures, one characterisation is that a structure $\mathbb{C}$ is diagonalisable if and only if the two projection homomorphisms $\pi_1, \pi_2 \colon \mathbb{C} \times \mathbb{C} \to \mathbb{C}$ (defined as $\pi_i(x_1, x_2) = x_i$) are connected by some sequence of homomorphisms $\psi \colon \mathbb{C} \times \mathbb{C} \to \mathbb{C}$ 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 [38]. For these and many other equivalent definitions of diagonalisability, cf. [37, Corollary 6.3 and Theorem 3.6] with $J = V(H)$.

Our main result for minimisation is an approximation scheme for instances $(\mathbb{A}, \mathbb{C})$ where $\mathbb{A}$ comes from a Baker class and $\mathbb{C}$ is diagonalisable.

Theorem I.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 $\mathrm{G}(\mathbb{A}) \in \mathcal{G}$ and $\mathbb{C}$ is diagonalisable, we can find a solution of value at most $(1+\varepsilon)\operatorname{minval}(\mathbb{A}, \mathbb{C})$ 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 $(\mathcal{O}(n^2), \mathcal{O}(n))$-efficiently Baker [22, Theorem 2.1], Theorem I.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).

Maximum Solution

For maximisation, we extend the results of [32], 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 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 $\mathrm{G}(\mathbb{A}) \in \mathcal{G}$ 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 is the following.

Theorem I.2. *Let $\mathcal{G}$ be a class of graphs that is fractionally-treewidth-fragile. Then 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 $\operatorname{maxval}(\mathbb{A}, \mathbb{C})$ and $(1+\varepsilon) \cdot \operatorname{maxval}(\mathbb{A}, \mathbb{C})$ in time $(|\mathbb{A}| + |\mathbb{C}|)^{k(\varepsilon)}$, where $k(\varepsilon)$ is a function depending on $\mathcal{G}$ only.

The algorithm in Theorem I.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 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 [32], we conjecture that Max-Sol$_{\mathcal{G}}$ does not admit a PTAS for any $\mathcal{G}$ that is not fractionally-treewidth-fragile. Since Max-Sol$_{\mathcal{G}}$ is strictly more general (by allowing negative infinite values), this conjecture might be easier to prove than the one in [32].

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 Max-Sol$_{\mathcal{G}}$ of value at least $(1-\varepsilon)\cdot \mathrm{maxval}(\mathbb{A},\mathbb{C})$ in time $|\mathbb{A}|\cdot|\mathbb{C}|^{k(\varepsilon)}$.

We complement Theorem I.2 with simple constructions which show that it is impossible to extend other results of [32] to the setting of crisp constraints. In [32] the notion of *pliability* is defined, which is a left-hand side restriction that takes the whole structure $\mathbb{A}$ into account, not only its Gaifman graph; this allowed the authors of [32] to show that the same framework applies not only to sparse, fractionally-treewidth-fragile instances, but also to dense structures. We define an analogous notion of *strong pliability* and show in Lemma IV.13, similarly to [32], the existence of a PTAS 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 I.2.) However, in the full version [35] 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 [35] that the Max-Sol problem is hard to approximate even when the left-hand side structures are restricted to cliques.

Paper organisation: Section II introduces basic notations and defines the studied computational problems. The main result for minimisation, Theorem I.1, is technical and proved in the full version [35]. In Section III, we present the main ideas in the special case of planar structures. The main result for maximisation, Theorem I.2, is proved in Section IV. Some of the proofs are deferred to the full version [35].

II. 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 $\mathrm{Set}(\mathbf{x})$ 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 $\mathrm{ar}(f)$. For a set of values $\Omega\subseteq\mathbb{Q}\cup\{-\infty,+\infty\}$, an *$\Omega$-valued structure* $\mathbb{A}$ over a signature σ (or σ-structure, for short) is a finite domain A together with a function $f^{\mathbb{A}}: A^{\mathrm{ar}(f)}\to\Omega$ 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 $\mathrm{tup}(\mathbb{A})$ to be the set of all pairs $(f,\mathbf{x})$ such that $f\in\sigma$ and $\mathbf{x}\in A^{\mathrm{ar}(f)}$; and $\mathrm{tup}_{>0}(\mathbb{A})$ to be the set of all pairs $(f,\mathbf{x})\in\mathrm{tup}(\mathbb{A})$ with $f^{\mathbb{A}}(\mathbf{x})>0$.

We consider the following computational problem.

Definition II.1. An instance of the *general-valued CSP* (VCSP) consists of an ordered pair of σ-structures $(\mathbb{A},\mathbb{C})$. For a mapping $h\colon A\to C$, we define the value of h to be

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

The goal is to find the minimum or maximum value over all possible mappings $h\colon A\to C$, denoted $\mathrm{minval}(\mathbb{A},\mathbb{C})$ or $\mathrm{maxval}(\mathbb{A},\mathbb{C})$, 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 $(f,\mathbf{x})\in\mathrm{tup}(\mathbb{A})$. 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 $\mathbf{x}\in X^{\mathrm{ar}(f)}$.

Following the influential work on decision CSPs by Grohe, Schwentick, and Segoufin [25], and Grohe [39], 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 $\mathrm{G}(\mathbb{A})$, 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 $\mathrm{p}(\mathbb{A}):=\mathrm{p}(\mathrm{G}(\mathbb{A}))$ to be the parameter of the Gaifman graph of $\mathbb{A}$. In particular, the *treewidth* of $\mathbb{A}$ is defined as $\mathrm{tw}(\mathbb{A}):=\mathrm{tw}(\mathrm{G}(\mathbb{A}))$. (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 [40].)

Relational structures: A *relational σ-structure* $\mathbb{C}$ includes for each symbol $f\in\sigma$ a relation $f^{\mathbb{C}}\subseteq C^{\mathrm{ar}(f)}$. We will view relational structures as $\{0,\infty\}$-valued structures by associating each function $f^{\mathbb{C}}: C^{\mathrm{ar}(f)}\to\{0,\infty\}$ 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 $\psi: C\to D$ that satisfies, for every $f\in\sigma$ and every $\mathbf{x}\in C^{\mathrm{ar}(f)}$, $f^{\mathbb{D}}(\psi(\mathbf{x}))\leq f^{\mathbb{C}}(\mathbf{x})$.

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

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

Maximum Solution

Definition II.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 II.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 $C^{\mathrm{ar}(f)}$. 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 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 $\mathrm{G}(\mathbb{A}) \in \mathcal{G}$ and $\mathbb{C}$ is a Max-Sol structure.

Observe that every $\mathbb{Q}_{\geq 0}$-valued structure is a Max-Sol structure; thus 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 [32].

Remark II.4. The "downward monotone Strict-CSP" from [3] corresponds to Definition II.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 $C^{\mathrm{ar}(f)}$ 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 II.5 (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 $\mathrm{ar}(f) > 1$ 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 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 $\mathrm{G}(\mathbb{A}) \in \mathcal{G}$ and $\mathbb{C}$ is a Min-Sol structure.

Remark II.6. The "upward monotone Strict-CSP" from [3] corresponds to Definition II.5 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 II.7. 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 doesn't 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 $\mathrm{minval}(\mathbb{A}, \mathbb{C}) = 0$. Otherwise, $\mathrm{minval}(\mathbb{A}, \mathbb{C}) \geq 1$. Note that $f^{\mathbb{C}}$ respects the partial order $\sqsubseteq_{c_\top}$, but it does not respect any total order on C.

III. Minimisation on planar structures

A. Diagonalisability

Briceño, Bulatov, Dalmau, and Larose defined the concepts of product structure, dismantlability, homomorphisms, adjacency, and link graph for relational structures [37]. 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 [37] (for relational structures) when viewed as $\{0, \infty\}$-valued structures.

The intuition is that, given two (valued) σ-structures $\mathbb{C}$ and $\mathbb{D}$, we call $\psi : C \to D$ a *homomorphism* from $\mathbb{C}$ to $\mathbb{D}$ if ψ is a homomorphism from $\mathrm{Feas}(\mathbb{C})$ to $\mathrm{Feas}(\mathbb{D})$ and from $\mathrm{Opt}(\mathbb{C})$ to $\mathrm{Opt}(\mathbb{D})$. It will be more convenient to consider both the $\mathrm{Feas}(\mathbb{C})$ and $\mathrm{Opt}(\mathbb{C})$ simultaneously. Thus with every structure $\mathbb{C}$ we will associate a relational structure $\mathrm{Rel}[\mathbb{C}]$, defined as follows.

Definition III.1. If $\mathbb{C}$ is a σ-structure, let $\sigma' = \bigcup_{f \in \sigma} \{f_1, f_2\}$ be the signature that for each $f \in \sigma$ includes the relational symbols f_1 of f_2 of the same arity as f. Define the relational σ'-structure $\mathrm{Rel}[\mathbb{C}]$ over the domain C as follows: for each $f \in \sigma$, let $f_1^{\mathrm{Rel}[\mathbb{C}]} = \mathrm{Feas}(f^{\mathbb{C}}) = \{\mathbf{x} \mid f^{\mathbb{C}}(\mathbf{x}) < \infty\}$ and $f_2^{\mathrm{Rel}[\mathbb{C}]} = \mathrm{Opt}(f^{\mathbb{C}}) = \{\mathbf{x} \mid f^{\mathbb{C}}(\mathbf{x}) = 0\}$.

We can now define the concepts of interest for structures $\mathbb{C}$ via the already existing concepts for relational structures $\mathrm{Rel}[\mathbb{C}]$ from [37]. We use the following observation.

Observation III.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 $\psi : C \to D$ is a *homomorphism* if ψ is a homomorphism from $\mathrm{Rel}[\mathbb{C}]$ to $\mathrm{Rel}[\mathbb{D}]$. Equivalently, ψ is a homomorphism if there exists $M > 0$ such that for all $(f, \mathbf{x}) \in \mathrm{tup}(\mathbb{C})$,

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

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 $\mathbb{C}$.

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

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

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 $\mathbb{C}_\ell$. A structure $\mathbb{C}$ is *diagonalisable* if $\mathbb{C}^2$ dismantles to the substructure induced by its diagonal $\Delta(C^2)=\{(c,c)\mid c\in C\}$.

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

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

Thus a dominated by b in $\mathbb{C}$ if and only if the function $s\colon C\to C\setminus\{a\}$ 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 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 [37, Theorem 3.6] for relational structures but the result easily extends to structures.

Theorem III.3. *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 $\psi: C^2\to C$ is idempotent if $\psi(c,c)=c$ for all $c\in C$.)

Proof. This follows from the fact that $\mathrm{Rel}[\mathbb{C}^2]=\mathrm{Rel}[\mathbb{C}]^2$ and that our definitions are the same as those of [37, Theorem 3.6] applied to $\mathbb{H}=\mathbb{C}$ and $J=\Delta(C)$. Specifically a function $\phi: C\to C$ is a homomorphism from $\mathbb{C}$ to $\mathbb{C}$ if and only if it is a homomorphism from $\mathrm{Rel}[\mathbb{C}]$ to $\mathrm{Rel}[\mathbb{C}]$. Similarly, a is dominated by b in $\mathbb{C}^2$ if and only if a is dominated by b in $\mathrm{Rel}[\mathbb{C}^2]=\mathrm{Rel}[\mathbb{C}]^2$. Thus $\mathbb{C}$ is diagonalisable if and only if $\mathrm{Rel}[\mathbb{C}]^2$ dismantles to its full diagonal (not just any subset of it). Further $\phi,\psi: C\to C$ are adjacent homomorphisms from $\mathbb{C}$ to $\mathbb{C}$ if and only if they are adjacent homomorphisms from $\mathrm{Rel}[\mathbb{C}]$ to $\mathrm{Rel}[\mathbb{C}]$. Finally, π_1,π_2 are connected by a path of adjacent idempotent homomorphisms if and only if they are J-connected by any homomorphisms in $L(\mathrm{Rel}[\mathbb{C}]^2,\mathbb{C})$ in the sense of [37]. $\square$

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

Lemma III.4. *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 $\phi: C^2\to\Delta(C^2)$ 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 $(f,(\mathbf{x},\mathbf{y}))\in\mathrm{tup}(\mathbb{C}^2)$ with $(x_i,y_i)=a$ and $n:=\mathrm{ar}(f)>1$, we have

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

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

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

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 $\mu: C^2\to C$ 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 [37] 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 $\mathrm{Rel}[\mathbb{C}]$ 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 III.3 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.

B. PTAS

Baker's approach relies on the following structural property of planar graphs, which is e.g. a direct consequence of [41, Theorem 83].

Lemma III.5. *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 ith layer of a BFS from v_0). Then, the subgraph induced by any t consecutive layers $G[L_i \cup L_{i+1} \cup \cdots \cup L_{i+t+1}]$ has treewidth at most $3t$.*

Theorem III.6. *Let $\mathcal{P}$ be the class of planar graphs. Then, for any $\varepsilon > 0$ and any VCSP instance $(\mathbb{A}, \mathbb{C})$ with $\mathrm{G}(\mathbb{A}) \in \mathcal{P}$ and $\mathbb{C}$ diagonalisable, we can find a solution of value at most $(1+\varepsilon)\,\mathrm{minval}(\mathbb{A}, \mathbb{C})$ 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 $\mathrm{val}_{\mathbb{B}}(h)$ for the value of an assignment $h : B \to C$ with respect to the instance $(\mathbb{B}, \mathbb{C})$, and write $\mathrm{val}(\cdot)$ for $\mathrm{val}_{\mathbb{A}}(\cdot)$ by default.

By Theorem III.3 there is a sequence of adjacent homomorphisms $\psi_1, \ldots, \psi_\ell$ 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}, $i \in \{1, \ldots, \ell - 1\}$. Let $k \coloneqq \lceil \frac{2M}{\varepsilon} \rceil$.

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

$$B_j^i \coloneqq L_{jk-i} \cup \cdots \cup L_{jk-i+k},$$

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 $O^1, \ldots, O^k$ are disjoint. Consider an optimal solution $h^* : A \to C$ for the VCSP instance $(\mathbb{A}, \mathbb{C})$. As the O^i are all disjoint, there exists $i^* \in [k]$ with

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

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 $\mathrm{val}(h')$.

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

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

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

For each j, the Gaifman graph $\mathrm{G}(\mathbb{A}^+[B_j])$ has treewidth at most $\mathcal{O}((k+1)\ell) = \mathcal{O}(M\ell/\varepsilon)$ by Lemma III.5. Thus for each j, we can find a treewidth decomposition [42] and compute an optimal solution h_j to $(\mathbb{A}^+[B_j], \mathbb{C})$ in total time $|\mathbb{A}| \cdot |\mathbb{C}|^{\mathcal{O}(M\ell/\varepsilon)}$ via standard dynamic programming. Then by optimality of h_j,

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

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:

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

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 $h' : A \to C$ be defined as follows

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

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

1) If $\mathrm{Set}(\mathbf{x}) \not\subseteq O$, then there is a unique j such that $\mathrm{Set}(\mathbf{x}) \subseteq B_j$, 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 $h'(x) = \psi_\ell\big(h_{j-1}(x), h_j(x)\big) = \pi_2\big(h_{j-1}(x), h_j(x)\big) = h_j(x)$, or analogously x_i is in the first layer of $B_j \cap B_{j+1}$ and $\psi_1 = \pi_1$. Thus
$$f^{\mathbb{C}}(h'(\mathbf{x})) = f^{\mathbb{C}}(h_j(\mathbf{x})). \tag{3}$$

2) Else, if $\mathrm{Set}(\mathbf{x}) \subseteq O$, then there is a unique j such that $\mathrm{Set}(\mathbf{x}) \subseteq B_j \cap B_{j+1} = L_{(j+1)k-i^*}$. Since $\mathrm{Set}(\mathbf{x})$ is contained in two consecutive layers, there is some s such

that all vertices in $\mathrm{Set}(\mathbf{x})$ are at distance $((j+1)k - i^*)\ell + s$ or $((j+1)k-i^*)\ell+s+1$ from v_0. Thus

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

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

$$\begin{aligned} f^{\mathbb{C}}(h'(\mathbf{x})) &\le M \cdot f^{\mathbb{C}^2}\big(h_j(\mathbf{x}), h_{j+1}(\mathbf{x})\big) \\ &= M \cdot \big(f^{\mathbb{C}}\big(h_j(\mathbf{x})\big) + f^{\mathbb{C}}(h_{j+1}(\mathbf{x}))\big)\,. \end{aligned} \tag{4}$$

Thus, by (2) to (4),

$$\begin{aligned} \mathrm{val}_{\mathbb{A}}(h') &\le \sum_j \mathrm{val}_{\mathbb{A}+[B_j]}(h_j) \\ &\le (1+\varepsilon)\,\mathrm{val}(h^*) = (1+\varepsilon)\,\mathrm{minval}(\mathbb{A},\mathbb{C}), \end{aligned}$$

and so h' is the solution we seek. □

Remark III.7. In Theorem III.6 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 [37, Theorem 3.6] (extended as in Theorem III.3) 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 I.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 III.4, we have the following corollary.

Corollary III.8. *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 $(1+\varepsilon)\,\mathrm{minval}(\mathbb{A},\mathbb{C})$ in time $|\mathbb{A}|\cdot c^{1/\varepsilon}$, where c depends on $\mathbb{C}$ only.*

We remark the proof yields $c^{1/\varepsilon} = |\mathbb{C}|^{\mathcal{O}(M\ell/\varepsilon)}$, and for Min-Sol structures Lemma III.4 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 $|\mathbb{C}|^{\mathcal{O}(1/\varepsilon)}$.

IV. 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 $\mathcal{O}(k)$, 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 [32]; 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.

A. Pliability

Definition IV.1. For two left-hand side σ-structures $\mathbb{A},\mathbb{B}$, we say $\mathbb{A}$ *strongly overcasts* $\mathbb{B}$, denoted $\mathbb{A} \succeq \mathbb{B}$, if for all Max-Sol σ-structures $\mathbb{C}$, $\mathrm{maxval}(\mathbb{A},\mathbb{C}) \ge \mathrm{maxval}(\mathbb{B},\mathbb{C})$.

In contrast, [32] defined (weak) overcasting in terms of $\mathbb{Q}_{\ge 0}$-valued structures $\mathbb{C}$ only, instead of the wider class of Max-Sol structures. The "strong" qualifier is only to avoid confusion with [32]: we will not consider weak overcasts in this paper, nor analogous weak variants of the definitions given below.

Definition IV.2. The *strong opt-distance* between two left-hand side σ-structures $\mathbb{A}$ and $\mathbb{B}$ is defined as

$$\mathrm{d}_{opt}(\mathbb{A},\mathbb{B}) \coloneqq \inf\{\varepsilon \mid \mathbb{A} \succeq e^{-\varepsilon}\mathbb{B} \text{ and } \mathbb{B} \succeq e^{-\varepsilon}\mathbb{A}\}.$$

Observation IV.3. Using the fact that $\mathrm{maxval}(\lambda\mathbb{A},\mathbb{C}) = \lambda\,\mathrm{maxval}(\mathbb{A},\mathbb{C})$, it is an easy exercise to see that $\mathrm{d}_{opt}(\mathbb{A},\mathbb{B}) = \infty$ if exactly one of $\mathrm{maxval}(\mathbb{A},\mathbb{C}), \mathrm{maxval}(\mathbb{B},\mathbb{C})$ is $-\infty$, or exactly one of them is 0, for some Max-Sol σ-structure $\mathbb{C}$; otherwise

$$\begin{aligned} \mathrm{d}_{opt}(\mathbb{A},\mathbb{B}) = \sup\{\varepsilon \mid \mathbb{A} \not\succeq e^{-\varepsilon}\mathbb{B} \text{ or } \mathbb{B} \not\succeq e^{-\varepsilon}\mathbb{A}\} = \\ \sup_{\mathbb{C}} |\ln \mathrm{maxval}(\mathbb{A},\mathbb{C}) - \ln \mathrm{maxval}(\mathbb{B},\mathbb{C})|. \end{aligned}$$

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 [32], one can prove that treedepth, or the Hadwiger number, give rise to equivalent definitions.

Definition IV.4. A class of $\mathbb{Q}_{\ge 0}$-valued structures $\mathcal{A}$ is *strongly* p-*pliable* with respect to a graph parameter p if for all $\varepsilon > 0$ there exists $k = k(\varepsilon)$ such that for every σ-structure $\mathbb{A} \in \mathcal{A}$ there exists a $\mathbb{Q}_{\ge 0}$-valued σ-structure $\mathbb{B}$ with $\mathrm{p}(\mathbb{B}) \le k$ and $\mathrm{d}_{opt}(\mathbb{A},\mathbb{B}) \le \varepsilon$.

B. Duality

Definition IV.5 (partial functions and homomorphisms)**.** For a partial function $g\colon A \to B$ and a tuple $\mathbf{x} \in A^n$, $g(\mathbf{x})$ is defined as $(g(x_1),\ldots,g(x_n)) \in B^n$ if all coordinates are defined, and is undefined otherwise. For $\mathbf{y} \in B^n$, we define $g^{-1}(\mathbf{x}) \coloneqq \{\mathbf{x} \in A^n \mid g(\mathbf{x}) \text{ is defined and equal to } \mathbf{y}\}$.

For left-hand side σ-structures $\mathbb{A},\mathbb{B}$, a *partial homomorphism* from $\mathbb{A}$ to $\mathbb{B}$ is a partial function $g\colon A \to B$ such that: for any positive tuple $(f,\mathbf{x}) \in \mathrm{tup}_{>0}(\mathbb{A})$, there is a positive tuple $(f,\mathbf{y}) \in \mathrm{tup}_{>0}(\mathbb{B})$ 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 $\mathrm{p\text{-}hom}(\mathbb{A},\mathbb{B})$.

Remark IV.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

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

In particular $f^{\mathbb{B}^+}(\mathbf{x}) = f^{\mathbb{B}}(\mathbf{x})$ for $\mathbf{x} \in B^n$. Let $\mathrm{Pos}(\mathbb{A})$ 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 $\mathrm{Pos}(\mathbb{A})$ to $\mathrm{Pos}(\mathbb{B}^+)$ (undefined assignments are the same as assignments to $\star$).

Lemma IV.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 $\mathbb{C}$, $\mathrm{maxval}(\mathbb{A}, \mathbb{C}) \geq \mathrm{maxval}(\mathbb{B}, \mathbb{C})$;*
- *there is a distribution of partial homomorphisms $\omega\colon$ p-hom$(\mathbb{A}, \mathbb{B}) \to \mathbb{Q}_{\geq 0}$ $(\sum_g \omega(g) = 1)$ such that*

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

(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 $\mathbf{x} \in A^{\mathrm{ar}(f)}$ such that $g(\mathbf{x})$ is defined and equal to $\mathbf{y}$.)

C. PTAS

We first define the Sherali-Adams LP hierarchy [43] for Max-Sol. Let $(\mathbb{A}, \mathbb{C})$ be an instance of Max-Sol over a signature σ and let $k \geq \max_{f \in \sigma} \mathrm{ar}(f)$. 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* [43] of $(\mathbb{A}, \mathbb{C})$ is the linear program given in Fig. 1, denoted by $\mathrm{SA}_k(\mathbb{A}, \mathbb{C})$, which has one variable $\lambda(X, s)$ for each $X \in \binom{A}{\leq k}$ and each $s\colon X \to C$. We denote by $\mathrm{maxval}_k(\mathbb{A}, \mathbb{B})$ the optimum value of $\mathrm{SA}_k(\mathbb{A}, \mathbb{C})$, and define $\mathrm{maxval}_k(\mathbb{A}, \mathbb{B}) = -\infty$ if $\mathrm{SA}_k(\mathbb{A}, \mathbb{C})$ is infeasible.

Observation IV.8. Let $(\mathbb{A}, \mathbb{C})$ be an instance of Max-Sol, $k \geq \max_{f \in \sigma} \mathrm{ar}(f)$ and $\lambda \geq 0$. Then, $\mathrm{maxval}(\lambda\mathbb{A}, \mathbb{C}) = \lambda \,\mathrm{maxval}(\mathbb{A}, \mathbb{C})$ and $\mathrm{maxval}_k(\lambda\mathbb{A}, \mathbb{C}) = \lambda \,\mathrm{maxval}_k(\mathbb{A}, \mathbb{C})$.

Observation IV.9. Let $(\mathbb{A}, \mathbb{C})$ be an instance of Max-Sol. Then, for any $k \geq \max_{f \in \sigma} \mathrm{ar}(f)$, $\mathrm{maxval}_k(\mathbb{A}, \mathbb{C}) \geq \mathrm{maxval}(\mathbb{A}, \mathbb{C})$.

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

The following easy result shows that an appropriate level of the Sherali-Adams relaxation is exact for bounded treewidth.

Proposition IV.10. *Let $(\mathbb{A}, \mathbb{C})$ be an instance of Max-Sol and $k \geq \mathrm{tw}(\mathbb{A})$. Then,* $\mathrm{maxval}_k(\mathbb{A}, \mathbb{C}) = \mathrm{maxval}(\mathbb{A}, \mathbb{C})$.

Definition IV.11. Let $\mathbb{A}$ and $\mathbb{B}$ be left-hand side σ-structures, and $k \geq \max_{f \in \sigma} \mathrm{ar}(f)$. We write $\mathbb{A} \succeq_k \mathbb{B}$ if for all Max-Sol σ-structures $\mathbb{C}$ we have $\mathrm{maxval}_k(\mathbb{A}, \mathbb{C}) \geq \mathrm{maxval}_k(\mathbb{B}, \mathbb{C})$.

Using the dual characterisation of strong overcasts (Lemma IV.7), we can show the following.

Proposition IV.12. *Let $\mathbb{A}$ and $\mathbb{B}$ be left-hand side σ-structures, and $k \geq \max_{f \in \sigma} \mathrm{ar}(f)$. 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 IV.13. *Let $\mathbb{A}$ be a left-hand side σ-structure, $\varepsilon \geq 0$ be small and $k \geq \max_{f \in \sigma} \mathrm{ar}(f)$. Suppose that there exists a left-hand side σ-structure $\mathbb{B}$ such that $\mathrm{d}_{opt}(\mathbb{A}, \mathbb{B}) \leq \varepsilon$ and $\mathrm{tw}(\mathbb{B}) \leq k$. Then, for every right-hand side σ-structure $\mathbb{C}$, we have that*

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

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

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

and so

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

by Proposition IV.12. From Observations IV.8 and IV.9 and Proposition IV.10 we obtain that,

$$\begin{aligned}\mathrm{maxval}(\mathbb{A}, \mathbb{C}) &\leq \mathrm{maxval}_k(\mathbb{A}, \mathbb{C}) \leq e^{\varepsilon} \,\mathrm{maxval}_k(\mathbb{B}, \mathbb{C}) \\ &= e^{\varepsilon} \,\mathrm{maxval}(\mathbb{B}, \mathbb{C}) \leq e^{2\varepsilon} \,\mathrm{maxval}(\mathbb{A}, \mathbb{C}).\end{aligned}$$

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

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

Corollary IV.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.*

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

D. Fragility and pliability

To give Dvořák's definition of fractional fragility [21] we first define ε-thin distributions.

Definition IV.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 $v \in V$.

Definition IV.16. For a graph parameter p and a number k, we define a *$(\mathrm{p} \leq k)$-modulator* of a graph G to be a set $X \subseteq V(G)$ such that $\mathrm{p}(G - X) \leq k$. A *fractional* $(\mathrm{p} \leq k)$*-modulator* is a distribution π of such modulators X. We say that a class of graphs $\mathcal{G}$ is *fractionally-*p*-fragile* if for every $\varepsilon > 0$ there is a k such that every $G \in \mathcal{G}$ has an ε-thin fractional $(\mathrm{p} \leq k)$-modulator.

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

$$
\begin{aligned}
\max \sum_{(f,\mathbf{x})\in \operatorname{tup}\mathbb{A},\ s\colon \operatorname{Set}(\mathbf{x})\to C} & \lambda(\operatorname{Set}(\mathbf{x}), s) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(s(\mathbf{x})) && \\
\lambda(X,s) = & \sum_{r\colon Y\to C,\, r|_X = s} \lambda(Y,r) && \text{for } X\subseteq Y \in \tbinom{A}{\le k} \text{ and } s\colon X\to C \\
\sum_{s\colon X\to C} \lambda(X,s) = & \ 1 && \text{for } X\in \tbinom{A}{\le k} \\
\lambda(\operatorname{Set}(\mathbf{x}), s) = & \ 0 && \forall (f,\mathbf{x})\in \operatorname{tup}(\mathbb{A}) \text{ with } f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{C}}(s(\mathbf{x})) = -\infty \\
\lambda(X,s) \ge & \ 0 && \text{for } X\in \tbinom{A}{\le k} \text{ and } s\colon X\to C
\end{aligned}
$$

Fig. 1. $\mathrm{SA}_k(\mathbb{A},\mathbb{C})$, the Sherali-Adams relaxation of level $k \ge \max_{f\in\sigma} \operatorname{ar}(f)$ of $(\mathbb{A},\mathbb{C})$.

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 IV.17. *Let* p *be a monotone*[2] *graph parameter such that* $\mathrm{p}(G \uplus H) \le \max(\mathrm{p}(G), \mathrm{p}(H))$ *for all graphs* G *and* H *and* $\mathrm{p}(G) \le \mathrm{p}(G - v) + 1$ *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. Then* $\mathcal{A}$ *is strongly* p*-pliable.*

The proof closely follows the proof of [44, 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 $(\mathrm{p} \le k)$-modulator such that for every $v \in V(G)$,

$$
\Pr_{X\sim\pi}[v \in X] \le \varepsilon. \tag{5}
$$

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 $g_X : A \to (A - X) \cup \{\star_X\}$ 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 $\mathrm{p}(\mathrm{G}(\mathbb{B}_X)) \le \mathrm{p}(\mathrm{G}(\mathbb{A}) - X) + 1 \le k + 1$, and so $\mathrm{p}(\mathrm{G}(\mathbb{B})) \le k + 1$.

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

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

[2] $\mathrm{p}(H) \le \mathrm{p}(G)$ for all graphs G and subgraphs H of G.

Define $g\colon B \to A$ 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 $g \in \text{p-hom}(\mathbb{B},\mathbb{A})$. Consider the overcast $\omega'\colon \mathbb{B} \to (1 - r\varepsilon)\mathbb{A}$ that is deterministically g. To check that ω' is indeed a strong overcast, let $(f, \mathbf{x}) \in \operatorname{tup}(\mathbb{A})$. Then $\mathbf{x}$ is covered by copies in $\mathbb{B}_X$ for those X that do not intersect $\mathbf{x}$, hence

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

where the final inequality follows by (5), the union bound, and the fact that $|\mathbf{x}| \le r$. Hence by Lemma IV.7 applied to ω and ω',

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

By construction $\mathrm{p}(\mathrm{G}(\mathbb{B})) \le k + 1$. As r is fixed and $\varepsilon > 0$ was arbitrary, this implies that $\mathcal{A}$ is strongly p-pliable. ☐

Proof of Theorem I.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 IV.17, we have that $\mathcal{A}$ is strongly tw-pliable. By Corollary IV.14, Max-Sol$_{\mathcal{G}}$ admits a PTAS. ☐

If we only look at Gaifman graphs, one cannot use the presented approach to go beyond fractionally-treewidth-fragile classes. This is because [44, Lemma 6.1] together with the above Lemma IV.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 $\mathbb{Q}_{\ge 0}$-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 the full version [35], 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.

REFERENCES

[1] S. Khanna and R. Motwani, "Towards a syntactic characterization of PTAS," in *Proc. 28th Annual ACM Symposium on the Theory of Computing (STOC'96)*. ACM, 1996, pp. 329–337.

[2] S. Khanna, M. Sudan, L. Trevisan, and D. P. Williamson, "The Approximability of Constraint Satisfaction Problems," *SIAM J. Comput.*, vol. 30, no. 6, pp. 1863–1920, 2001.

[3] A. Kumar, R. Manokaran, M. Tulsiani, and N. K. Vishnoi, "On LP-based approximability for strict CSPs," in *Proc. 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'11)*. SIAM, 2011, pp. 1560–1573.

[4] P. Jonsson, F. Kuivinen, and G. Nordh, "MAX ONES generalized to larger domains," *SIAM J. Comput.*, vol. 38, no. 1, pp. 329–365, 2008.

[5] G. Gutin, P. Hell, A. Rafiey, and A. Yeo, "A dichotomy for minimum cost graph homomorphisms," *European J. Combin.*, vol. 29, no. 4, pp. 900–911, 2008.

[6] R. Takhanov, "A Dichotomy Theorem for the General Minimum Cost Homomorphism Problem," in *Proc. 27th International Symposium on Theoretical Aspects of Computer Science (STACS'10)*, ser. LIPIcs. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2010, pp. 657–668.

[7] P. Hell and A. Rafiey, "The Dichotomy of Minimum Cost Homomorphism Problems for Digraphs," *SIAM J. Discrete Math*, vol. 26, no. 4, pp. 1597–1608, 2012.

[8] P. Hell, M. Mastrolilli, M. M. Nevisi, and A. Rafiey, "Approximation of Minimum Cost Homomorphisms," in *Proc. 20th Annual European Symposium on Algorithms (ESA12)*, ser. Lecture Notes in Computer Science, vol. 7501. Springer, 2012, pp. 587–598.

[9] A. Rafiey, A. Rafiey, and T. Santos, "Toward a Dichotomy for Approximation of H-Coloring," in *Proc. 46th International Colloquium on Automata, Languages, and Programming (ICALP'19)*, ser. LIPIcs, vol. 132. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019, pp. 91:1–91:16.

[10] P. Jonsson and G. Nordh, "Introduction to the maximum solution problem," in *Complexity of Constraints: An Overview of Current Research Themes*, ser. Lecture Notes in Computer Science, N. Creignou, P. G. Kolaitis, and H. Vollmer, Eds., vol. 5250. Springer, 2008, pp. 255–282.

[11] P. Jeavons, A. A. Krokhin, and S. Živný, "The complexity of valued constraint satisfaction," *Bull. EATCS*, vol. 113, 2014. [Online]. Available: http://eatcs.org/beatcs/index.php/beatcs/article/view/266

[12] K. Makarychev and Y. Makarychev, "Approximation algorithms for csps," in *The Constraint Satisfaction Problem: Complexity and Approximability*, ser. Dagstuhl Follow-Ups, A. A. Krokhin and S. Živný, Eds. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017, vol. 7, pp. 287–325.

[13] B. S. Baker, "Approximation Algorithms for NP-Complete Problems on Planar Graphs," *J. ACM*, vol. 41, no. 1, pp. 153–180, 1994.

[14] M. Grohe, "Local tree-width, excluded minors, and approximation algorithms," *Combinatorica*, vol. 23, no. 4, pp. 613–632, 2003.

[15] M. DeVos, G. Ding, B. Oporowski, D. P. Sanders, B. Reed, P. Seymour, and D. Vertigan, "Excluding any graph as a minor allows a low tree-width 2-coloring," *J. Comb. Theory B*, vol. 91, no. 1, pp. 25–41, 2004.

[16] E. D. Demaine, M. T. Hajiaghayi, and K. Kawarabayashi, "Algorithmic Graph Minor Theory: Decomposition, Approximation, and Coloring," in *Proc. 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. IEEE Computer Society, 2005, pp. 637–646.

[17] M. V. Marathe, H. B. Hunt III, and R. E. Stearns, "Level-treewidth property, exact algorithms and approximation schemes," in *Proc. 29th Annual ACM Symposium on Theory of Computing (STOC'97)*. ACM, 6 1997. [Online]. Available: https://www.osti.gov/biblio/471394

[18] A. Grigoriev and H. L. Bodlaender, "Algorithms for graphs embeddable with few crossings per edge," *Algorithmica*, vol. 49, no. 1, pp. 1–11, 9 2007.

[19] H. B. Hunt III, M. V. Marathe, V. Radhakrishnan, S. Ravi, D. J. Rosenkrantz, and R. E. Stearns, "NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs," *J. Algorithms*, vol. 26, no. 2, pp. 238–274, 1998.

[20] P. Berman and M. Karpinski, "On some tighter inapproximability results (extended abstract)," in *Proc. 26th International Colloquium on Automata, Languages and Programming (ICALP'99)*, ser. Lecture Notes in Computer Science, vol. 1644. Springer, 1999, pp. 200–209.

[21] Z. Dvořák, "Sublinear separators, fragility and subexponential expansion," *European J. Combin.*, vol. 52, pp. 103–119, 2016.

[22] ——, "Baker game and polynomial-time approximation schemes," in *Proc. 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'20)*. SIAM, 2020, pp. 2227–2240.

[23] ——, 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.

[24] D. P. Dailey, "Uniqueness of colorability and colorability of planar 4-regular graphs are NP-complete," *Discrete Math.*, vol. 30, no. 3, pp. 289–293, 1980.

[25] M. Grohe, T. Schwentick, and L. Segoufin, "When is the evaluation of conjunctive queries tractable?" in *Proc. 33rd Annual ACM Symposium on Theory of Computing (STOC'01)*. ACM, 2001, pp. 657–666.

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

[27] P. Jonsson, A. A. Krokhin, and F. Kuivinen, "Hard constraint satisfaction problems have hard gaps at location 1," *Theor. Comput. Sci.*, vol. 410, no. 38-40, pp. 3856–3874, 2009.

[28] P. Fulla, H. Uppman, and S. Živný, "The complexity of Boolean surjective general-valued CSPs," *ACM Trans. Comput. Theory*, vol. 11, no. 1, pp. 4:1–4:31, 2019.

[29] M. C. Cooper and S. Živný, "Hybrid tractable classes of constraint problems," in *The Constraint Satisfaction Problem: Complexity and Approximability*, ser. Dagstuhl Follow-Ups, A. A. Krokhin and S. Živný, Eds. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017, vol. 7, pp. 113–135.

[30] C. Carbonnel, M. Romero, and S. Živný, "The complexity of general-valued CSPs seen from the other side," in *Proc. 59th IEEE Annual Symposium on Foundations of Computer Science (FOCS'18)*. IEEE Computer Society, 2018, pp. 236–246.

[31] V. Kolmogorov, A. A. Krokhin, and M. Rolínek, "The complexity of general-valued csps," *SIAM J. Comput.*, vol. 46, no. 3, pp. 1087–1110, 2017.

[32] M. Romero, M. Wrochna, and S. Živný, "Treewidth-Pliability and PTAS for Max-CSPs," in *Proc. 2021 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'21)*. SIAM, 2021, pp. 473–483.

[33] S. Khot, "On the power of unique 2-prover 1-round games," in *Proc. 34th Annual ACM Symposium on Theory of Computing (STOC'02)*. ACM, 2002, pp. 767–775.

[34] P. Raghavendra, "Optimal algorithms and inapproximability results for every CSP?" in *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC'08)*. ACM, 2008, pp. 245–254.

[35] B. F. Mezei, M. Wrochna, and S. Živný, "PTAS for sparse general-valued CSPs," Tech. Rep., 2020, arXiv:2012.12607.

[36] G. R. Brightwell and P. Winkler, "Gibbs measures and dismantlable graphs," *J. Comb. Theory, Ser. B*, vol. 78, no. 1, pp. 141–166, 2000.

[37] R. Briceño, A. A. Bulatov, V. Dalmau, and B. Larose, "Dismantlability, connectedness, and mixing in relational structures," in *Proc. 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, ser. LIPIcs, C. Baier, I. Chatzigiannakis, P. Flocchini, and S. Leonardi, Eds., vol. 132. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019, pp. 29:1–29:15.

[38] A. A. Bulatov, A. A. Krokhin, and B. Larose, "Dualities for constraint satisfaction problems," in *Complexity of Constraints: An Overview of Current Research Themes*, ser. Lecture Notes in Computer Science, N. Creignou, P. G. Kolaitis, and H. Vollmer, Eds., vol. 5250. Springer, 2008, pp. 93–124.

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

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

[41] H. L. Bodlaender, "A partial k-arboretum of graphs with bounded treewidth," vol. 209, pp. 1–45, 1998.

[42] ——, "A linear time algorithm for finding tree-decompositions of small treewidth," in *Proc. 25th Annual ACM Symposium on Theory of Computing (STOC'93)*. ACM, 1993, p. 226–234.

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

[44] M. Romero, M. Wrochna, and S. Živný, "Treewidth-Pliability and PTAS for Max-CSPs," Tech. Rep., 10 2020.