

Pliability and Approximating Max-CSPs*

Miguel Romero
Pontificia Universidad Católica de Chile, Chile
mgromero@uc.cl

Marcin Wrochna
University of Warsaw, Poland
m.wrochna@mimuw.edu.pl

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

September 19, 2023

Abstract

We identify a sufficient condition, *treewidth-pliability*, that gives a polynomial-time algorithm for an arbitrarily good approximation of the optimal value in a large class of Max-2-CSPs parameterised by the class of allowed constraint graphs (with arbitrary constraints on an unbounded alphabet). Our result applies more generally to the maximum homomorphism problem between two rational-valued structures.

The condition unifies the two main approaches for designing a polynomial-time approximation scheme. One is Baker’s layering technique, which applies to sparse graphs such as planar or excluded-minor graphs. The other is based on Szemerédi’s regularity lemma and applies to dense graphs. We extend the applicability of both techniques to new classes of Max-CSPs. On the other hand, we prove that the condition cannot be used to find solutions (as opposed to approximating the optimal value) in general.

Treewidth-pliability turns out to be a robust notion that can be defined in several equivalent ways, including characterisations via size, treedepth, or the Hadwiger number. We show connections to the notions of fractional-treewidth-fragility from structural graph theory, hyperfiniteness from the area of property testing, and regularity partitions from the theory of dense graph limits. These may be of independent interest. In particular we show that a monotone class of graphs is hyperfinite if and only if it is fractionally-treewidth-fragile and has bounded degree.

1 Introduction

The problem of finding a *maximum cut* in a graph (Max-Cut) is one of the most studied problems from Karp’s original list of 21 NP-complete problems [56]. While Max-Cut is NP-hard

*An extended abstract of this work appeared in the Proceedings of SODA’21 [80]. Work done while Miguel Romero and Marcin Wrochna were at the University of Oxford. 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. This work was also supported by UKRI EP/X024431/1. For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission. All data is provided in full in the results section of this paper.

to solve optimally, there is a trivial 0.5-approximation algorithm [81] and the celebrated 0.878-approximation algorithm of Goemans and Williamson [42]. Papadimitriou and Yannakakis established that Max-Cut is Max-SNP-hard [77]. By the work of Arora, Lund, Motwani, Sudan, and Szegedy [5] this implies that, unless $P=NP$, there is no polynomial-time approximation scheme (PTAS) for Max-Cut in general graphs. However, non-trivial results exist for important special cases. On the one hand, Max-Cut is solvable exactly in planar graphs, as shown by Hadlock [50], and more generally, Max-Cut admits a PTAS on graph classes excluding a fixed minor, as shown by Demaine, Hajiaghayi, and Kawarabayashi [23]. On the other hand, Arora, Karger, and Karpinski showed a PTAS for Max-Cut in dense graphs [4], where a graph class is dense if every graph in it contains at least a constant fraction of all possible edges.

Max-Cut is an example of *maximum constraint satisfaction problem* (Max-CSP), although a very special one (the alphabet size is 2, in particular constant, and every constraint uses the same symmetric predicate “ $x \neq y$ ” of arity 2). Another well-known example is Max- r -SAT, with alphabet size 2 and r -ary clauses. Motivated by results on planar, excluded-minor, and dense graph classes, our goal in this paper is to understand the following question:

What structure allows for the existence of a PTAS for Max-CSPs?

We adopt a permissive definition of PTAS here: given a Max-CSP instance and an arbitrarily small $\varepsilon > 0$, the goal is to find a $(1 - \varepsilon)$ multiplicative approximation of the value of an optimal solution to the instance (but, unlike in most papers, we do not require that the algorithm should find a solution achieving the bound).

We focus on two computational problems. First, we study the general Max-2-CSP(\mathcal{G}) problem parameterised by the class of underlying constraint graphs (a.k.a. *primal* or *Gaifman* graphs). The input is a graph $G \in \mathcal{G}$, an alphabet Σ_v for each vertex, and a valued constraint $f_{uv}: \Sigma_u \times \Sigma_v \rightarrow \mathbb{Q}_{\geq 0}$ for each edge uv . The goal is to find an assignment $h(v) \in \Sigma_v$ maximising $\sum_{uv} f_{uv}(h(u), h(v))$. Similarly, in Max- r -CSP(\mathcal{G}) a constraint may appear on any r -clique in G . The constraints are arbitrary (non-negative) and the alphabets are not fixed, making the problem very expressive.¹

Second, we consider a more general framework called the *maximum homomorphism problem* (Max-Hom) of computing the maximum value of any map between two given $\mathbb{Q}_{\geq 0}$ -valued structures \mathbb{A} and \mathbb{B} ; the value will be denoted by $\text{opt}(\mathbb{A}, \mathbb{B})$ (see Section 2 for precise definitions). Intuitively, the left-hand-side structure describes the (weighted) scopes of the constraints and the right-hand-side structure describes the different types of constraints. Following Grohe’s notation [47], for a class of structures \mathcal{A} we denote by $\text{Max-Hom}(\mathcal{A}, -)$ the restriction of Max-Hom to instances (\mathbb{A}, \mathbb{B}) with $\mathbb{A} \in \mathcal{A}$ and \mathbb{B} arbitrary. This framework captures the Max- r -CSP(\mathcal{G}) problem as a particular case: it is equivalent to $\text{Max-Hom}(\mathcal{A}_{\mathcal{G}}^{(r)}, -)$, where by $\mathcal{A}_{\mathcal{G}}^{(r)}$ we denote the class of all valued structures with an underlying graph in \mathcal{G} and arity r . Another example is the case of *graph* Max-CSP, by which we mean a Max-2-CSP that uses the same symmetric predicate in all constraints (as in Max-Cut or Max- q -Cut); this case is equivalent to $\text{Max-Hom}(\mathcal{A}, -)$ where the structures in \mathcal{A} are graphs.

¹One could attempt to generalise counting problems by maximising $\prod_{uv} f_{uv}(h(u), h(v))$ instead, or equivalently its logarithm $\sum_{uv} \log f_{uv}(h(u), h(v))$. However, the requirement $f_{uv} \geq 0$ and the approximation ratio change. This changes the complexity: for example, approximating the number of 3-colourings requires deciding whether there is at least one in polynomial time, which is NP-hard already in 4-regular planar graphs [15].

The question of what structure allows to solve Max-CSPs *exactly* in polynomial time is well understood. A standard dynamic approach works for $\text{Max-}r\text{-CSP}(\mathcal{G})$ when \mathcal{G} is a class of graphs of bounded treewidth. Grohe, Schwentick, and Segoufin [48] in fact proved the converse: if \mathcal{G} has unbounded treewidth then $\text{Max-}r\text{-CSP}(\mathcal{G})$, in fact already deciding the existence of a solution satisfying all constraints, cannot be solved in polynomial time (assuming $\text{FPT} \neq \text{W}[1]$). Grohe’s theorem [47] then extended it to the more general framework: for a class of relational (or $\{0, 1\}$ -valued) structures \mathcal{A} of bounded arity, the decision problem $\text{Hom}(\mathcal{A}, -)$ can be solved in polynomial time if and only if the *cores* of structures in \mathcal{A} have bounded treewidth. (The core is the smallest homomorphically equivalent substructure; for example, bipartite graphs all have the single edge graph K_2 as a core, so $\text{Hom}(\mathcal{A}, -)$ is easy when \mathcal{A} is a class of bipartite graphs). This was recently extended further to exact optimisation with valued structures by Carbonnel, Romero, and Živný [11].

Max- r -CSPs do not admit a PTAS in general, since already Max-Cut does not. On the other hand, the techniques that give PTASes for Max-Cut on sparse and dense graphs apply more generally (in fact to a variety of problems beyond Max-CSPs). Our main contribution is a unifying condition, *treewidth-pliability*, that captures all known PTASes for $\text{Max-}r\text{-CSP}(\mathcal{G})$ and $\text{Max-Hom}(\mathcal{A}, -)$ problems.

We call a class of structures \mathcal{A} *tw-pliable* if it is uniformly close to structures of bounded treewidth. More formally, for any $\varepsilon > 0$ there is a $k = k(\varepsilon)$ such that every structure in \mathcal{A} has an ε -close structure with treewidth at most k . Here we consider two structures \mathbb{A} and \mathbb{B} to be ε -close if $\text{opt}(\mathbb{A}, \mathbb{C})$ is ε -close to $\text{opt}(\mathbb{B}, \mathbb{C})$ for all \mathbb{C} (details in Section 2.3; this notion of distance, which we also characterise combinatorially, may be of independent interest). While the structure of bounded treewidth is not known and cannot be efficiently computed, we show that the Sherali-Adams LP relaxation gives a PTAS for $\text{Max-Hom}(\mathcal{A}, -)$.

Theorem 1.1. *If \mathcal{A} is a tw-pliable class of structures of bounded arity, then $\text{Max-Hom}(\mathcal{A}, -)$ admits a PTAS.*

We emphasise the generality of Theorem 1.1.² Firstly, the computational problem (Max-Hom) captures many fundamental problems, including graph homomorphisms [52], Max-Cut, Max-DiCut, Max-SAT, Max-CSPs, and query related problems coming from database theory [47]. Secondly, the notion of pliability captures many previously discovered cases of structures that admit a PTAS. In particular, we now discuss how Theorem 1.1 extends the applicability of the two main approaches for obtaining PTASes.

1.1 Sparse structures: Baker’s technique and fragility

Perhaps the best known technique for solving problems on planar graphs is Lipton and Tarjan’s planar separator theorem [63] and the divide & conquer approach it enables [64]. It can be used to give a PTAS for Max-CSPs with fixed alphabet size on planar graphs (this extends to excluded-minor graphs [2] and more [35]) of bounded degree.

This approach was superseded by Baker’s technique [6], which provides better running times and is easily applied to general Max- r -CSPs on arbitrary planar graphs (see e.g. [57]). The idea is very elegant: we partition a planar graph into Breadth-First-Search layers, remove every ℓ -th layer, and show that the remaining components of $\ell - 1$ consecutive layers have bounded treewidth (and so can be solved exactly). By trying different starting layers we can

²However, the generality comes at a cost, as detailed in Section 1.3: while an approximate optimum can be found, an approximate solution cannot be constructed unless $\text{P} = \text{NP}$.

ensure that the removed layers intersect an unknown optimal solution at most $\mathcal{O}(\frac{1}{\ell})$ times, giving a $1 \pm \mathcal{O}(\frac{1}{\ell})$ approximation.

From planar graphs this was extended to graphs of bounded genus by Eppstein [40] and later to all graph classes excluding a fixed minor by Grohe [46] and Demaine et al. [23]. The structural property needed for this approach, originally proved for excluded-minor graphs by DeVos, Ding, Oporowski, Sanders, Reed, Seymour, and Vertigan [24], is *tw-fragility*: they can be partitioned into any constant number of parts such that removing any one part leaves a graph of bounded treewidth. As shown by Hunt, Marathe, and Stearns [70, 54] (see also [55]) as well as Grigoriev and Bodlaender [45], the same property applies to some geometrically-defined graph classes that do not exclude any minor. One example is intersection graphs of unit disks whose centers are at least some constant apart (capturing some applications of the closely related shifting technique of Hochbaum and Maass [53] for geometric packing and covering problems). Another example is 1-planar graphs, or more generally graphs drawn on a fixed surface with a bounded number of intersections per edge.

An important generalisation, *fractional-tw-fragility*, was introduced by Dvořák [29]: it suffices that the parts whose removal results in a graph of bounded treewidth are nearly-disjoint (Definition 4.2). This applies to d -dimensional variants of the geometric classes mentioned above (for any constant d), in particular to d -dimensional grids, which are not tw-fragile [8]; this also includes classes of polynomial growth [60, 31]. Another large family of fractionally-tw-fragile classes are classes of bounded degree with strongly sublinear separators [29] (equivalently, bounded degree and polynomial expansion [35]). For such concrete examples of fragile classes, known proofs show that the nearly-disjoint parts can be computed efficiently. A PTAS can then easily be designed from the definition [29].

We show that the assumption about efficient construction is not needed. We do this by proving that if \mathcal{G} is *any* fractionally-tw-fragile class of graphs (intuitively, any class where a Baker-like technique is known to work), then the class $\mathcal{A}_{\mathcal{G}}^{(r)}$ of all possible structures of bounded arity r and with Gaifman graph in \mathcal{G} is tw-pliable.

Theorem 1.2. *Let \mathcal{G} be a fractionally-tw-fragile class of graphs. Then $\mathcal{A}_{\mathcal{G}}^{(r)}$ is tw-pliable for every r . Consequently, $\text{Max-}r\text{-CSP}(\mathcal{G})$ admits a PTAS.*

This captures all graph classes \mathcal{G} where a PTAS for $\text{Max-}r\text{-CSP}(\mathcal{G})$ is known.

1.2 Dense structures: the regularity lemma

It is perhaps more surprising that dense structures admit a PTAS. Here a class is *dense* if a constant factor of all possible constraints is present in every structure in the class, e.g. graphs with $\Omega(n^2)$ edges. Arora, Karger, and Karpinski [4] showed that $\text{Max-}r\text{-CSP}$ s admit a PTAS in the dense regime if the alphabet size is constant (in fact Boolean); de la Vega [18] independently gave a PTAS for dense Max-Cut. Frieze and Kannan [41] proved that these results are essentially possible because of Szemerédi’s regularity lemma [85]: intuitively, every graph can be approximated to within an additive $\pm \varepsilon n^2$ error by a random graph (with a constant number of parts, depending on ε only, so that the edges between two parts form a uniformly random graph of some density). For dense graphs, the additive error translates to a relative error, giving a PTAS. They also showed a variant of the regularity lemma that is still applicable to $\text{Max-}r\text{-CSP}$ s with constant alphabet size, yet avoids its infamous tower-type dependency on ε .

Goldreich, Goldwasser, and Ron [44] connected these results to the area of *property testing*, spawning an entirely new direction of research. They gave constant-time algorithms estimating the optimum value of some graph Max-CSPs. In fact, Alon, de la Vega, Kannan, and Karpinski [1] (see also Andersson and Engebretsen [3]) showed that Max- r -CSPs with a fixed alphabet can be approximated with accuracy $\pm \varepsilon n^r$ by sampling a constant number of vertices (polynomial in $\frac{1}{\varepsilon}$) and finding the optimum on the resulting (constant-size) induced substructure.

None of these results apply to any Max- r -CSP(\mathcal{G}) and Max-Hom($\mathcal{A}, -$) problem, that is, to unbounded alphabets. We give the first such example: undirected graphs with $\Omega(n^2)$ edges.

Theorem 1.3. *Let $c > 0$ and let \mathcal{A} be a class of graphs with at least cn^2 edges. Then \mathcal{A} is tw-pliable. Consequently, Max-Hom($\mathcal{A}, -$) admits a PTAS.*

(Note here the graphs in \mathcal{A} are input structures, not just Gaifman graphs of input structures). We also show that this cannot be extended to general CSPs: already for the class of tournaments—that is, orientations of complete graphs—a PTAS is impossible, (cf. Corollary 8.5 in Section 8) and indeed, this class is not tw-pliable (cf. Remark 7.4).

1.3 Robustness of pliability

The notion of treewidth-pliability not only unifies the different existing algorithmic techniques but it is also quite robust: treewidth-pliability captures a valued analogue of “homomorphic equivalence” (e.g. bipartite graphs, or 3-colourable graphs where each edge is contained in exactly one triangle, cf. Examples 2.14 and 2.15 in Section 2.5) as well as small edits: if \mathcal{A} is a pliable class of graphs, say, then the class of graphs obtained by adding or removing $o(m)$ edges from m -edge graphs in \mathcal{A} is again pliable (Lemma 2.13 in Section 2.5). However, this generality comes at a price. First, we show that even for fixed alphabet size, although the approximate optimum value can be found, an approximate solution cannot be constructed (unless $P = NP$, cf. Example 2.16 in Section 2.5). Second, unlike in some of the previous results for more restricted classes, our result does not give an EPTAS (i.e., with the degree of the polynomial time bound independent of ε) for fixed alphabet size (cf. Question 9.2). Finally, the use of strong versions of the regularity lemma yields tower-type dependencies on the approximation ratio ε in the dense case.

In the definition of treewidth-pliability we approximate structures by comparing their $\text{opt}()$ values and we ask them to be close to structures where the problem can be solved exactly. This is a non-constructive and very general definition. In fact, it is not inconceivable that this captures all tractable cases, i.e., that Max-Hom($\mathcal{A}, -$) has a PTAS if and only if \mathcal{A} is tw-pliable. Nevertheless, we show a variety of equivalent combinatorial definitions, which allow us to place a fairly tight bound on what pliability is, structurally.

For classes of the form $\mathcal{A}_{\mathcal{G}}^{(r)}$, that is, if we only restrict the underlying Gaifman graphs, we show that pliability collapses to fractional fragility. In this sense we understand the “sparse” setting exactly.

Lemma 1.4. *Let \mathcal{G} be a class of graphs. The following are equivalent, for any $r \geq 2$:*

- \mathcal{G} is fractionally-tw-fragile;
- $\mathcal{A}_{\mathcal{G}}^{(r)}$ is tw-pliable.

In general, we can replace treewidth with other parameters of the Gaifman graph: size (number of vertices), treedepth, denoted by td , Hadwiger number (maximum clique minor size), or maximum connected component size, which we denote by cc .

Theorem 1.5. *Let \mathcal{A} be any class of structures. The following are equivalent:*

- \mathcal{A} is td -pliable;
- \mathcal{A} is tw -pliable;
- \mathcal{A} is Hadwiger-pliable.

If structures in \mathcal{A} have bounded signatures, then the following are equivalent to the above as well:

- \mathcal{A} is size-pliable;
- \mathcal{A} is cc -pliable.

Classes of structures with bounded signatures (see Section 2 for precise definitions) correspond to Max-CSP instances with a bounded number of constraint types; e.g. maximum graph homomorphism. For example, any class of dense graphs as in Theorem 1.3 is in fact size-pliable. An example of a class with unbounded signatures is any class of the form $\mathcal{A}_{\mathcal{G}}^{(r)}$ (we do not consider infinite signatures, but there are arbitrarily many symbols in those signatures). Theorem 1.5 allows us to give concrete and general examples of classes that are *not* tw -pliable: the class of orientations of graphs in \mathcal{G} , where \mathcal{G} is *any* class of unbounded average degree (Lemma 2.19 in Section 2.6), or any class of 3-regular graphs with unbounded girth (Lemma 2.25 in Section 2.6).

Finally, as a side result, we connect *hyperfiniteness* to fragility. A class of graphs \mathcal{G} is called *hyperfiniteness* if for every $\varepsilon > 0$ there is a $k = k(\varepsilon)$ such that in every $G \in \mathcal{G}$ one can remove an at-most- ε fraction of edges to obtain a graph with connected components of size at most k . For a monotone class of graphs (closed under taking subgraphs), hyperfiniteness easily implies bounded degree. It is an important notion in property testing: many results in sparse graphs were generalised by the statement that every property of hyperfinite graphs is testable [76]. The idea, originating in the work of Benjamini, Schramm, and Shapira [7] and Hassidim, Kelner, Nguyen, and Onak [51], is that following the approach of Lipton and Tarjan, graphs with sufficiently sublinear separators, such as planar or excluded-minor graphs [2], can be recursively partitioned into bounded-size components, which for bounded-degree graphs gives hyperfiniteness (see e.g. [14, Cor. 3.2] for a slightly stronger property, cf. [72]). This allows, analogously as in the dense case, to give a constant-size approximate description of such graphs by sampling constant-radius balls in them [76]. See [43] for a book on property testing and [61] for a recent improvement for excluded-minor graphs.

We show that a monotone class \mathcal{G} is hyperfinite if and only if it is fractionally- tw -fragile and has bounded degree. In fact, replacing the parameter treewidth by the maximum size of a connected component in a graph, we have:

Theorem 1.6. *Let \mathcal{G} be a monotone class of graphs. The following are equivalent:*

- \mathcal{G} is hyperfinite;
- \mathcal{G} is fractionally- tw -fragile and has bounded degree;
- \mathcal{G} is fractionally- cc -fragile;
- $\mathcal{A}_{\mathcal{G}}^{(r)}$ is cc -pliable (for any $r \geq 2$).

The equivalence of the second and third bullet points was shown by Dvořák [29, Observation 15, Corollary 20], while for the third and fourth the proof is established by a generalisation of Lemma 1.4, cf. Lemma 4.7.

Hyperfiniteness originates from the study of amenable groups and graphs limits, with motivations in geometry and mathematical physics [38]. The unexpected connection with fractional fragility already found an application in that area: Elek [39] showed that our Theorem 1.6 gives that last missing implication in proving the equivalence of some properties of infinite, bounded-degree graphs (in particular “uniform local amenability” and “Property A”), which answers a question of Brodzki, Niblo, Spakula, Willett and Wright [10].

1.4 Related work

While this paper focuses on Max- r -CSPs, Baker’s technique and the regularity lemma apply to many more problems. In fact Khanna and Motwani [57] argued that most known PTAS algorithms can be derived from three canonical optimisation problems on planar graphs, the first being Max-CSP and the latter two being so-called Max-Ones and Min-Ones CSPs (also solvable with Baker’s technique). One of the very few results that did not fit their framework was the PTAS for dense Max-Cut. A follow-up work by Mezei, Wrochna, and Živný [71] on the extended abstract of this work [80] extended some of the results of the present paper to Min- and Max-CSPs with crisp constraints, which include the Max-Ones and Min-Ones CSPs mentioned above.

Generic frameworks extending Baker’s technique include the bidimensionality theory of Demaine, Fomin, Hajiaghayi, and Thilikos [21] and its application in the design of PTASes by Demaine and Hajiaghayi [22] (which is however limited to minor-closed graph classes); monotone FO problems on minor-closed graph classes by Dawar, Grohe, Kreutzer, and Schweikardt [17]; and the idea of Baker games, introduced by Dvořák [30] (see also [34]). The latter gives conditions stronger than fractional-tw-fragility, but useful for problems beyond Max-CSPs, and achievable for all examples known to be fractionally fragile. The work of Dvořák and Lahiri [33], which appeared after the present paper, gives a PTAS on fractionally-tw-fragile classes of graphs for problems incomparable with Max-CSPs, namely monotone maximisation problems expressible in terms of distances.

De la Vega and Karpinski [19, 20] extended the dense approach to subdense cases ($\Omega(\frac{n^2}{\log n})$ edges) for specific problems such as MaxCut and Max-2-SAT. In contrast, they show that Max-Cut on graphs with $\Omega(n^{2-\delta})$ edges is hard to approximate, for any $\delta > 0$.

The best known approximation algorithm for general Max-2-CSPs is due to Charikar, Hajiaghayi, and Karloff [13] and achieves an approximation factor of $\mathcal{O}((nq)^{1/3})$, where n is the number of variables and q is the alphabet size. On the hardness side, Dinur, Fischer, Kindler, Raz, and Safra [27] showed that $\mathcal{O}(2^{\log^{1-\delta}(nq)})$ -approximation of Max-2-CSPs is NP-hard. Manurangsi and Moshkovitz [68] gave approximation algorithms for *dense* Max-2-CSPs with large alphabet size (but not PTASes). Manurangsi and Raghavendra [69] establish a tight trade-off between running time and approximation ratio for dense Max- r -CSPs for $r > 2$.

CSPs have also been extensively studied for fixed constraint types, i.e., Max-Hom($-, \mathbb{B}$) problems for fixed \mathbb{B} . Raghavendra showed that the best approximation ratio is always achieved by the basic SDP relaxation [78], assuming Khot’s unique games conjecture [58]. The exactly solvable cases were characterised by Thapper and Živný [86]. The approximation factor of graph Max-CSPs was studied by Langberg, Rabani, and Swamy [62].

1.5 Overview

In Section 2, we give formal definitions and present our basic tool: two structures \mathbb{A}, \mathbb{B} have similar values of $\text{opt}(-, \mathbb{C})$ if and only if there is a certain fractional cover, which we call an *overcast*, from \mathbb{A} to \mathbb{B} and from \mathbb{B} to \mathbb{A} . Section 2 then relates our notion of pliability with two notions of distances, and gives examples and non-examples of pliable classes of structures.

To prove that treewidth-pliability leads to a PTAS (Theorem 1.1) the main idea is that an overcast allows to show that the values of $\text{opt}(-, \mathbb{C})$ are still similar when we look at linear programming relaxations. The details, as well as the definition of the Sherali-Adams linear programming relaxation, are given in Section 3. In Section 4, we introduce equivalent definitions of fractional fragility and study their properties. This will allow us to prove Theorem 1.2 by showing how the definition implies suitable overcasts. This also allows us to establish Lemma 1.4. Theorem 1.5 is proved in Section 5. Theorem 1.6 on hyperfiniteness is proved in Section 6. Section 7 gives a proof of Theorem 1.3 on dense graphs.

We conclude with open questions in Section 9.

2 Preliminaries

2.1 Structures

A *signature* is a finite set σ of (function) symbols f , each with a specified finite arity $\text{ar}(f)$. We denote by $|\sigma|$ the number of symbols in the signature σ . A *structure* \mathbb{A} over a signature σ (or σ -structure \mathbb{A} , for short) is a finite domain A together with a function $f^{\mathbb{A}} : A^{\text{ar}(f)} \rightarrow \mathbb{Q}_{\geq 0}$ for each symbol $f \in \sigma$. We say that a class of structures has *bounded signatures* if for the signatures σ of structures in the class, $|\sigma|$ is bounded by a constant (so unbounded means arbitrarily many symbols; we do not consider infinite signatures). Note that a class of σ -structures (that is, structures over a fixed signature σ) has bounded signatures and bounded arities (the maximum arity occurring in σ is a finite constant).

We denote by A, B, C, \dots the domains of structures $\mathbb{A}, \mathbb{B}, \mathbb{C}, \dots$. For sets A and B , we denote by B^A the set of all mappings from A to B . 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 by $\text{tup}(\mathbb{A})_{>0}$ the pairs $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$ with $f^{\mathbb{A}}(\mathbf{x}) > 0$.

We denote $\|\mathbb{A}\|_{\infty} := \max_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x})$ and $\|\mathbb{A}\|_1 := \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x})$. 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 $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$.

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 $r \geq 2$ and a class of graphs \mathcal{G} , we denote by $\mathcal{A}_{\mathcal{G}}^{(r)}$ the class of σ -structures \mathbb{A} with $G(\mathbb{A}) \in \mathcal{G}$ and $\text{ar}(f) \leq r$ for every $f \in \sigma$. (Note that $\mathcal{A}_{\mathcal{G}}^{(r)}$ contains structures over distinct signatures.)

The *maximum homomorphism problem* (Max-Hom) is the following computational problem. An instance of Max-Hom consists of two structures \mathbb{A} and \mathbb{B} over the same signature. For a mapping $h : A \rightarrow B$, we define $\text{value}(h) = \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(h(\mathbf{x}))$. The goal is to find the maximum value over all possible mappings $h : A \rightarrow B$.³ We denote this value

³While called maximum homomorphism, we note that the maximisation is over all possible maps, not only homomorphisms, i.e., those that map non-zero tuples into non-zero tuples.

by $\text{opt}(\mathbb{A}, \mathbb{B})$. Note that when seen as a Max-CSP instance, the domain of the left-hand side structure \mathbb{A} is the variable set, while the domain of the right-hand side structure \mathbb{B} is the *alphabet*.

Example 2.1. Let $\sigma = \{f\}$ be a signature consisting of a single symbol f of arity $\text{ar}(f) = 2$. Let \mathbb{B} be a σ -structure with the domain $B = \{0, 1\}$ and let $f^{\mathbb{B}} : B^2 \rightarrow \mathbb{Q}_{\geq 0}$ be defined by $f^{\mathbb{B}}(x, y) = 1$ if $x \neq y$ and $f^{\mathbb{B}}(x, y) = 0$ if $x = y$. Given an undirected graph $G = (V, E)$, we can encode it as a σ -structure \mathbb{A} with the domain $A = V$ and with $f^{\mathbb{A}}(x, y) = 1$ if $\{x, y\} \in E$ and $f^{\mathbb{A}}(x, y) = 0$ otherwise. Now, the instance (\mathbb{A}, \mathbb{B}) of Max-Hom is the same as the Max-Cut problem in G . The Max-DiCut problem (in a directed graph (V, E)) would be cast as Max-Hom very similarly. The only differences would be in the definition of $f^{\mathbb{A}}$ and $f^{\mathbb{B}}$: $f^{\mathbb{A}}(x, y) = 1$ if $(x, y) \in E$ and 0 otherwise, $f^{\mathbb{B}}(x, y) = 1$ if $x = 0$ and $y = 1$ and 0 otherwise.

Example 2.2. An example of a problem that is not a Max-Hom is the Maximum Independent Set problem. Intuitively, the “no edges” constraints imposed on an independent set are strict. This problem can be, however, cast as an instance of a Max-Hom with both rational and (negative) infinite costs, cf. [71] for follow-up work.

Given a class \mathcal{A} of structures, $\text{Max-Hom}(\mathcal{A}, -)$ is the problem restricted to instances (\mathbb{A}, \mathbb{B}) of Max-Hom with $\mathbb{A} \in \mathcal{A}$ (it is a promise problem: algorithms are allowed to do anything when $\mathbb{A} \notin \mathcal{A}$). Recall that for a class of graphs \mathcal{G} , the problem $\text{Max-}r\text{-CSP}(\mathcal{G})$ is equivalent to $\text{Max-Hom}(\mathcal{A}_{\mathcal{G}}^{(r)}, -)$.⁴

2.2 Overcasts

Before we define pliability formally, it is useful to consider the following relation. The starting point of all our results is the equivalence of this relation to a more combinatorial notion: the existence of a certain fractional cover, which we shall call an *overcast*.

Definition 2.3. Let \mathbb{A} and \mathbb{B} be σ -structures. We say that \mathbb{A} *overcasts* \mathbb{B} , denoted $\mathbb{A} \succeq \mathbb{B}$ if, for all σ -structures \mathbb{C} , we have that $\text{opt}(\mathbb{A}, \mathbb{C}) \geq \text{opt}(\mathbb{B}, \mathbb{C})$.

A distribution over a finite set U is a function $\pi : U \rightarrow \mathbb{Q}_{\geq 0}$ such that $\sum_{x \in U} \pi(x) = 1$. We write $\mathbb{E}_{x \sim \pi} f(x)$ for $\sum_{x \in U} \pi(x) \cdot f(x)$ and $\Pr_{x \sim \pi}[\phi(x)]$ for $\mathbb{E}_{x \sim \pi}[\phi(x)]$, where $[\phi(x)]$ is 1 if x satisfies the predicate ϕ and 0 otherwise.

Given a map $g : A \rightarrow B$ and a tuple $\mathbf{x} = (x_1, \dots, x_m) \in A^m$, we write $g(\mathbf{x})$ for $(g(x_1), \dots, g(x_m))$; i.e., we apply g componentwise on \mathbf{x} . Hence, $g^{-1}(\mathbf{y}) = \{\mathbf{x} \mid g(\mathbf{x}) = \mathbf{y}\}$.

Definition 2.4. Let \mathbb{A} and \mathbb{B} be σ -structures. An *overcast* from \mathbb{A} to \mathbb{B} is a distribution ω over B^A such that for each $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})$ we have that

$$\mathbb{E}_{g \sim \omega} f^{\mathbb{A}}(g^{-1}(\mathbf{x})) \geq f^{\mathbb{B}}(\mathbf{x}).$$

Here $f^{\mathbb{A}}(g^{-1}(\mathbf{x}))$ denotes the sum of $f^{\mathbb{A}}(\mathbf{y})$ over $\mathbf{y} \in g^{-1}(\mathbf{x}) \subseteq A^{\text{ar}(f)}$.

⁴Note that $\text{Max-Hom}(\mathcal{A}_{\mathcal{G}}^{(r)}, -)$ is different from the maximum graph homomorphism problem $\text{Max-Hom}(\mathcal{G}, -)$. Indeed, graphs are also structures over the signature $\{e\}$ with one symbol of arity 2 (where $e^G(u, v) = [uv \text{ is an edge of } G]$, if the graph is not weighted). To avoid confusion, we use \mathcal{G} for a class of Gaifman graphs of some structures and \mathcal{A} for a class of graphs that are themselves used as input structures.

Intuitively, an overcast from \mathbb{A} to \mathbb{B} is a random function from A to B such that for each edge (tuple) in \mathbb{B} , its preimage has larger expected weight (value). In other words, each edge must be covered by at least its own weight, in expectation. The following is a consequence of Farkas' lemma, as shown in Appendix A.⁵

Proposition 2.5. $\mathbb{A} \succeq \mathbb{B}$ if and only if there is an overcast from \mathbb{A} to \mathbb{B} .

2.3 Pliability and graph parameters

Our definition of pliability involves a notion of distance which may be of independent interest. It quantifies the *relative* difference between two structures (as measured from the right by weighted multicut densities, in the language of Lovász's book on graph limits [66, Ch. 12]).

Definition 2.6. The *opt-distance* between two structures with the same signature is defined as:

$$d_{\text{opt}}(\mathbb{A}, \mathbb{B}) := \sup_{\mathbb{C}} |\ln \text{opt}(\mathbb{A}, \mathbb{C}) - \ln \text{opt}(\mathbb{B}, \mathbb{C})|.$$

Here $\ln 0 = -\infty$ and $|\ln 0 - \ln 0| = 0$. Equivalently, we can compare rescaled structures; by definition of \succeq and the fact that $\text{opt}(\lambda\mathbb{A}, \mathbb{C}) = \lambda \text{opt}(\mathbb{A}, \mathbb{C})$, we have:

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

One may think of $e^{\pm\varepsilon}$ as close to $1 \pm \varepsilon$. Formally $1 - \varepsilon \leq e^{-\varepsilon} \leq \frac{1}{1+\varepsilon} = 1 - \varepsilon + \mathcal{O}(\varepsilon^2)$ for $\varepsilon \geq 0$. Note that the first definition readily implies that opt-distance satisfies the triangle inequality (it defines a pseudometric).

Finally, a class is treewidth-pliable if it is uniformly close to structures of bounded treewidth:

Definition 2.7. A class of structures \mathcal{A} is *p-pliable* with respect to a graph parameter p if for every $\varepsilon > 0$, there is $k = k(\varepsilon)$ such that for every σ -structure $\mathbb{A} \in \mathcal{A}$ there is a σ -structure \mathbb{B} with $p(\mathbb{B}) \leq k$ and $d_{\text{opt}}(\mathbb{A}, \mathbb{B}) \leq \varepsilon$.

Thus to show tw-pliability of various classes, we will construct overcasts from structures \mathbb{A} in the class to $(1 - \varepsilon)\mathbb{B}$, for some \mathbb{B} of bounded treewidth, and from \mathbb{B} back to $(1 - \varepsilon)\mathbb{A}$.

Given a graph G , we will consider pliability for the following graph parameters:

- $\text{size}(G) = |V(G)|$ – the number of vertices of G ,
- $\text{cc}(G)$ – the maximum size of a connected component of G ,
- treedepth $\text{td}(G)$, which is a parameter due to Nešetřil and Ossona de Mendez [74], whose definition we recall below,
- treewidth $\text{tw}(G)$,⁶

⁵The definitions of the \succeq relation and of an overcast are analogous to the “improvement” relation and “inverse fractional homomorphisms” from [11]. Here, however, $\text{opt}()$ is maximising, not minimising, so inequalities in definitions are swapped. This has consequences such as the fact that mappings in the support of an overcast are in general not homomorphisms (mapping non-zero tuples to non-zero tuples), unlike for inverse fractional homomorphisms. The proof of Proposition 2.5 nevertheless is identical to the proof of [11, Proposition 3.6].

⁶We refer to [25] for the standard definitions of treewidth, pathwidth and minors.

- and finally the Hadwiger number $\text{Hadwiger}(G)$, which is the maximum k such that K_k is a minor of G . We use the Hadwiger number as an example of a broader “sparsity” parameter: for example, all planar graphs have Hadwiger number at most 4 (yet unbounded treewidth).

Definition 2.8. The *treedepth* $\text{td}(G)$ of a graph G is defined recursively as:

- $\max_i \text{td}(G_i)$, if G is disconnected with components G_i ;
- $\min_{v \in V(G)} \text{td}(G - v) + 1$, if G is connected and has more than one vertex;
- 1, if G has one vertex.

An equivalent definition is as follows: a *treedepth decomposition* of a graph G is a rooted forest T (a disjoint union of rooted trees) with $V(T) = V(G)$ such that for each $uv \in E(G)$, u is an ancestor or descendant of v in T . In other words, G is a subgraph of the transitive closure of a forest T directed towards roots. The treedepth of G is equal to the minimum depth among all such decompositions of G . Treedepth is a rather strict parameter: for example, stars have treedepth 2, but paths already have unbounded treedepth. In fact, a short proof shows that a class of graphs has bounded treedepth if and only if the length of the longest path is bounded [74].

Bounded size implies bounded cc implies bounded td implies bounded pathwidth (pw) implies bounded tw implies bounded Hadwiger number, more precisely:

$$\text{Hadwiger}(G) - 1 \leq \text{tw}(G) \leq \text{pw}(G) \leq \text{td}(G) - 1 \text{ and } \text{td}(G) \leq \text{cc}(G) \leq \text{size}(G).$$

Moreover, we also have the following inequality (less useful, because of the dependency on G):

$$\text{td}(G) \leq (\text{tw}(G) + 1) \cdot \log_2 |G|.$$

All these parameters are *monotone*, that is, $p(H) \leq p(G)$ for a subgraph H of a graph G . Their boundedness implies bounded average degree $\frac{2|E(G)|}{|V(G)|}$. More precisely, $\frac{2|E(G)|}{|V(G)|} \leq 2 \text{tw}(G)$ (because a graph of treewidth k has a vertex of degree at most k); Kostochka [59] proved $\frac{2|E(G)|}{|V(G)|} \leq \mathcal{O}(h\sqrt{h})$ where $h = \text{Hadwiger}(G)$.

The size of a structure \mathbb{A} is the number of vertices of its Gaifman graph: $\text{size}(\mathbb{A}) = |V(G(\mathbb{A}))|$. The other graph parameters are also defined in terms of the same parameter on the Gaifman graph of the structure; e.g., the treewidth of a structure \mathbb{A} is the treewidth of its Gaifman graph: $\text{tw}(\mathbb{A}) = \text{tw}(G(\mathbb{A}))$. In particular, since the edges of $G(\mathbb{A})$ come only from tuples in \mathbb{A} of non-zero weight, a rescaled structure $\lambda\mathbb{A}$, for $\lambda > 0$, has $G(\lambda\mathbb{A}) = G(\mathbb{A})$ so the parameters we consider do not change by rescaling; for $\lambda = 0$, $G(\lambda\mathbb{A})$ has no edges.

We will often prove the easy directions in various characterisations via the following observation that follows from the definition of pliability:

Observation 2.9. *If two graph parameters p and p' satisfy $p \leq p'$ then p' -pliability implies p -pliability.*

2.4 Opt-distance zero and edit-distance

In this section we define our relative version of edit distance and prove it upper-bounds opt-distance. We define the *edit distance* $d_1(\mathbb{A}, \mathbb{B})$ between two valued σ -structures \mathbb{A}, \mathbb{B} to be

$$d_1(\mathbb{A}, \mathbb{B}) := \min_{\text{bij. } \phi: A \rightarrow B} \sum_{f \in \sigma} \frac{\sum_{\mathbf{x} \in A^{\text{ar}(f)}} |f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{B}}(\phi(\mathbf{x}))|}{\min(\|\mathbb{A}_f\|_1, \|\mathbb{B}_f\|_1)}.$$

Here \mathbb{A}_f denotes the structure \mathbb{A} limited to the signature $\{f\}$, so $\|\mathbb{A}_f\|_1$ denotes $\sum_{\mathbf{x} \in A^{\text{ar}(f)}} f^{\mathbb{A}}(\mathbf{x})$. The following generalises the notion of “looplessness” in graphs.

Definition 2.10. A σ -structure \mathbb{A} is *loopless* if no tuple has a repetition. That is, for $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$ with $f^{\mathbb{A}}(\mathbf{x}) > 0$, \mathbf{x} consists of $\text{ar}(f)$ different elements of A .⁷

Lemma 2.11. *For loopless structures, the opt-distance is bounded linearly by the edit distance:*

$$d_{\text{opt}} \leq C_\sigma \cdot d_1,$$

where $C_\sigma = \max_{f \in \sigma} \text{ar}(f)^{\text{ar}(f)}$.

Proof. Let $d_1 = d_1(\mathbb{A}, \mathbb{B})$ and let $\phi: A \rightarrow B$ be a bijection minimising the expression in its definition. We will show that $e^{C_\sigma \cdot d_1} \mathbb{A} \succeq (1 + C_\sigma \cdot d_1) \mathbb{A} \succeq \mathbb{B}$. Symmetrically, $e^{C_\sigma \cdot d_1} \mathbb{B} \succeq \mathbb{A}$, hence $d_{\text{opt}} \leq C_\sigma \cdot d_1$, which will conclude our claim.

Observe that for $f \in \sigma$

$$\sum_{\mathbf{x} \in A^{\text{ar}(f)}} |f^{\mathbb{A}}(\mathbf{x}) - f^{\mathbb{B}}(\phi(\mathbf{x}))| = \sum_{\mathbf{x} \in B^{\text{ar}(f)}} |f^{\mathbb{B}}(\mathbf{x}) - f^{\mathbb{A}}(\phi^{-1}(\mathbf{x}))|.$$

Let $\delta := \frac{C_\sigma \cdot d_1}{1 + C_\sigma \cdot d_1}$, so $1 - \delta = \frac{1}{1 + C_\sigma \cdot d_1}$. To show $(1 + C_\sigma \cdot d_1) \mathbb{A} \succeq \mathbb{B}$, we construct an overcast ω from \mathbb{A} to $(1 - \delta) \mathbb{B}$ as follows. With probability $(1 - \delta)$ we map \mathbb{A} to \mathbb{B} with ϕ ; with probability δ we choose a tuple $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})$ at random with probability proportional to its contribution in d_1 , that is, $\frac{|f^{\mathbb{B}}(\mathbf{x}) - f^{\mathbb{A}}(\phi^{-1}(\mathbf{x}))|}{\min(\|\mathbb{A}_f\|_1, \|\mathbb{B}_f\|_1)} \cdot \frac{1}{d_1}$, and we map all of \mathbb{A} uniformly at random to the elements of this tuple. That is, after choosing $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})$, each tuple of \mathbb{A}_f gets mapped into \mathbf{x} with probability $\frac{1}{\text{ar}(f)^{\text{ar}(f)}}$ (assuming \mathbb{A} is loopless). Therefore, for each $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})$:

$$\begin{aligned} \sum_{g \in B^A} \omega(g) \cdot f^{\mathbb{A}}(g^{-1}(\mathbf{x})) &\geq (1 - \delta) \cdot f^{\mathbb{A}}(\phi^{-1}(\mathbf{x})) + \delta \cdot \frac{|f^{\mathbb{B}}(\mathbf{x}) - f^{\mathbb{A}}(\phi^{-1}(\mathbf{x}))|}{\min(\|\mathbb{A}_f\|_1, \|\mathbb{B}_f\|_1)} \cdot \frac{1}{d_1} \cdot \frac{\|\mathbb{A}_f\|_1}{\text{ar}(f)^{\text{ar}(f)}} \\ &\geq (1 - \delta) f^{\mathbb{A}}(\phi^{-1}(\mathbf{x})) + \frac{\delta}{C_\sigma \cdot d_1} \cdot |f^{\mathbb{B}}(\mathbf{x}) - f^{\mathbb{A}}(\phi^{-1}(\mathbf{x}))| \geq (1 - \delta) f^{\mathbb{B}}(\mathbf{x}), \end{aligned}$$

where the last inequality follows from $\frac{\delta}{C_\sigma \cdot d_1} = \frac{1}{1 + C_\sigma \cdot d_1} = 1 - \delta$. This shows that ω is indeed an overcast that certifies $\mathbb{A} \succeq (1 - \delta) \mathbb{B}$. \square

Observation 2.12. *Let \mathcal{A} be a tw-pliable class. Let \mathcal{B} be a class of structures such that for every $\mathbb{B} \in \mathcal{B}$ there is an $\mathbb{A} \in \mathcal{A}$ with $d_{\text{opt}}(\mathbb{B}, \mathbb{A}) \leq f(\text{tw}(\mathbb{B}))$, for some function $f(n) \xrightarrow{n \rightarrow \infty} 0$. Then \mathcal{B} is tw-pliable.*

⁷Equivalently, the Gaifman graph of \mathbb{A} is loopless.

Proof. Since \mathcal{A} is tw-pliable, for every $\varepsilon > 0$ there is a $k = k(\varepsilon)$ such that every structure \mathbb{A} in \mathcal{A} is ε -close to some structure of $\text{tw} \leq k(\varepsilon)$. To show that \mathcal{B} is tw-pliable, consider any $\varepsilon > 0$. Let n_ε be large enough so that $f(n) \leq \frac{\varepsilon}{2}$ for $n \geq n_\varepsilon$. Then for $\mathbb{B} \in \mathcal{B}$, either $\text{tw}(\mathbb{B}) \leq n_\varepsilon$ or \mathbb{B} is $f(\text{tw}(\mathbb{B})) \leq \frac{\varepsilon}{2}$ -close to some structure $\mathbb{A} \in \mathcal{A}$, which in turn is $\frac{\varepsilon}{2}$ -close to some structure of treewidth at most $k(\frac{\varepsilon}{2})$. In either case \mathbb{B} is ε -close to a structure of treewidth at most $\max(n_\varepsilon, k(\frac{\varepsilon}{2}))$. \square

2.5 Pliable examples

We give simple observations and examples: classes that are sufficiently close to pliable classes (in edit or opt-distance) are themselves pliable. We first consider simple examples with a fixed signature: graphs.

Lemma 2.13. *Let \mathcal{A} be a tw-pliable class of graphs. Let \mathcal{B} be a class of graphs such that every $H \in \mathcal{B}$ can be obtained from some $G \in \mathcal{A}$ by adding or removing $f(|E(H)|)$ edges, for some function $f(m) \in o(m)$. Then \mathcal{B} is tw-pliable.*

Proof. By Lemma 2.11, $d_{\text{opt}}(H, G) \leq 4d_1(H, G) \leq \frac{f(|E(H)|)}{|E(H)| - f(|E(H)|)} = f'(|E(H)|)$ for some function $f'(n) \xrightarrow{n \rightarrow \infty} 0$. This function can be upper-bounded by a monotonic function f'' decreasing to 0, say $f''(x) := \sup_{n \geq x} f'(n)$. Since $|E(H)| \geq \text{tw}(H)$, we conclude $d_{\text{opt}}(H, G) \leq f''(|E(H)|) \leq f''(\text{tw}(H))$. The claim follows by Observation 2.12. \square

Other simple examples arise from considering structures at opt-distance zero. This is a valued analogue of being homomorphically equivalent (see also *valued cores* in [11]).

Example 2.14. For every non-empty bipartite graph G , $d_{\text{opt}}(G, \lambda K_2) = 0$, for $\lambda = |E(G)|$. Therefore, since $\{\lambda K_2 : \lambda \in \mathbb{Q}_{\geq 0}\}$ is trivially tw-pliable, every class of bipartite graphs is tw-pliable.

Proof. A bipartite graph G admits a homomorphism h to K_2 . This gives an overcast showing $G \succeq \lambda K_2$: always map everything according to h . Conversely, mapping λK_2 uniformly at random to edges of G gives an overcast showing $\lambda K_2 \succeq G$. \square

Example 2.15. Let G be a 3-colourable graph such that every edge of G occurs in exactly one triangle. Then $d_{\text{opt}}(G, \lambda K_3) = 0$ for $\lambda = |E(G)|/3$. Hence the class of all such graphs is tw-pliable.

Proof. A 3-colouring of G corresponds to a homomorphism h to K_3 . Composing h with a random rotation of K_3 gives an overcast from G to λK_3 . Conversely, mapping λK_3 to a uniformly random triangle in G covers each edge with probability $\frac{1}{\lambda}$, giving an overcast from λK_3 to G . \square

The above idea also implies that our results cannot be extended to *finding* solutions. This is analogous to the hardness of finding a 3-colouring of a graph that is homomorphically equivalent to K_3 .

Example 2.16. There is a class of weighted graphs \mathcal{A} that is tw-pliable, yet for some $\varepsilon > 0$, there is no poly-time algorithm that finds a map $h: V(G) \rightarrow V(K_3)$ with $\text{value}(h) \geq (1 - \varepsilon)\text{opt}(G, K_3)$ for $G \in \mathcal{A}$, unless $P = NP$.

Proof. Let \mathcal{A} be the class of weighted graphs G satisfying $d_{\text{opt}}(G, \lambda K_3) = 0$ for $\lambda = \frac{\|G\|_1}{3}$. Suppose there is an algorithm as above for each $\varepsilon > 0$. There are constants ε_0, d such that it is NP-hard to distinguish 3-colourable graphs of maximum degree d from graphs where any map $h: V(G) \rightarrow V(K_3)$ miscolours more than an ε_0 fraction of edges [49]. We use our algorithm to solve the problem. Given an instance G , let G' be the weighted graph obtained by gluing a new triangle to every edge, and then assigning to every edge $e \in E(G')$ a weight $w(e)$ equal to the number of triangles it occurs in. Note $1 \leq w(e) \leq d$. Observe that if G was 3-colourable, then G' would be as well, hence $\text{opt}(G', K_3) = \|G'\|_1$. Moreover, G' would be in \mathcal{A} by an argument similar as in Example 2.15. Hence running the algorithm on G' , we would find a 3-colouring h which miscolours at most $\varepsilon \text{opt}(G', K_3) = \varepsilon \|G'\|_1$ of the total weight $\|G'\|_1$. Since $w \geq 1$, it miscolours at most $\varepsilon \|G'\|_1$ edges. Since $w \leq d$, this is at most $\varepsilon d |E(G')| \leq 3\varepsilon d |E(G)|$. Hence running the algorithm for $\varepsilon = \frac{\varepsilon_0}{3d}$ would find a colouring of the original graph G that miscolours at most $\varepsilon_0 |E(G)|$ edges. Therefore if we run this procedure for any G (regardless of its 3-colourability), then it either outputs a colouring as above, or we can conclude that G is not 3-colourable. \square

2.6 Non-pliable examples

In this section we give examples of non-pliable classes. In the process we show equivalent definitions of pliability (Lemmas 2.22 and 2.23).

We will use the following bound (in this section and in the proof of Lemma 8.3 in Section 8):

$$\sum_{i=0}^{\lambda m} \binom{m}{i} \leq 2^{H(\lambda)m},$$

where $H(\lambda)$ is a function which satisfies $\lim_{\lambda \rightarrow 0} H(\lambda) = 0$; specifically, the binary entropy function $H(\lambda) = \lambda \log_2(\frac{1}{\lambda}) - (1 - \lambda) \log_2(1 - \lambda)$.

Recall that by Lemma 1.4, for a class of Gaifman graphs \mathcal{G} , the class $\mathcal{A}_{\mathcal{G}}^{(2)}$ of all structures over \mathcal{G} is tw-pliable if and only if \mathcal{G} is fractionally-tw-fragile. So the simplest examples of non-pliable classes are $\mathcal{A}_{\mathcal{G}}^{(2)}$ for some non-fractionally-fragile \mathcal{G} . Fractional-tw-fragility implies bounded expansion (a notion from the theory of sparse graphs introduced by Nešetřil and Ossona de Mendez [75]) and sublinear separators, e.g., 3-regular expander graphs are not fractionally-tw-fragile, see [29]. Hence for \mathcal{G} the class of all 3-regular graphs, $\mathcal{A}_{\mathcal{G}}^{(2)}$ is not tw-pliable.

A somewhat more direct proof is to consider any class of 3-regular graphs of high girth. Thomassen [87] showed that such graphs behave much like graphs of high average degree. We use essentially the same proof below:

Lemma 2.17. *For $\delta > 0$ and $g \in \mathbb{N}$, every graph with average degree $\geq 2 + \delta$ and girth $\geq 3g$ has a minor with average degree $\geq g\delta + 2$.*

Proof. Let $G = (V, E)$. Without loss of generality assume that G is connected (otherwise take the component with the largest average degree). Let A_1, \dots, A_m be a partition of V into parts of size $|A_i| \geq g$ that induce connected subgraphs, with m maximum among such partitions (clearly one exists with $m = 1$).

We claim that each set A_i induces a tree. Indeed, consider any spanning tree T of $G[A_i]$ and let e be an edge of $G[A_i]$ outside of T . Then $T + e$ contains a unique cycle, which must

have length $\geq 3g \geq 2g$. Hence one can remove e and some other edge from this cycle to split it into two intervals with $\geq g$ vertices. Removing these two edges from $T + e$ splits it into two components spanning A_i with $\geq g$ vertices each. Hence A_i could be replaced with the vertex sets of these two components, contradicting the choice of m .

Similarly, we claim that every two sets A_i, A_j are connected by at most one edge. Otherwise two such edges together with spanning trees of A_i and A_j would form a unicyclic graph, which could be split as above into three connected parts with $\geq g$ vertices each.

Let $G' = (E', V')$ be the graph resulting by contracting the sets A_i . Since we contract sets of $\geq g$ vertices, $|V'| = m \leq \frac{|V|}{g}$. Since no two edges get identified and no loop gets created/removed in the process, the number of contractions is equal to $|E| - |E'|$ and to $|V| - |V'|$. Hence $|E'| = |E| - |V| + |V'| \geq (\frac{2+\delta}{2} - 1)|V| + |V'| \geq (g\frac{\delta}{2} + 1)|V'|$, so G' has average degree $\geq g\delta + 2$. (We note that each $G[A_i]$ had diameter $< 2g - 1$, as otherwise it could be split into two parts; hence the minor we obtain is relatively *shallow*). \square

Proposition 2.18. *Let $\delta > 0$ and let \mathcal{G} be a class of graphs with unbounded girth and average degree $\geq 2 + \delta$. Then \mathcal{G} is not fractionally-tw-fragile.*

Proof. Suppose that \mathcal{G} is fractionally-tw-fragile. Then, by Lemma 4.4, for $\varepsilon = \frac{\delta}{2(2+\delta)}$ there is a $k = k(\varepsilon)$ such that every graph in \mathcal{G} has a subset $F \subseteq E(G)$ with $|F| \leq \varepsilon|E(G)|$ such that $\text{tw}(G - F) \leq k$. Let $G \in \mathcal{G}$ be a graph with girth $\geq \frac{12k}{\delta}$. Let F be as above. Then $\text{tw}(G - F) \leq k$ and $2|E(G - F)| \geq (1 - \varepsilon) \cdot 2|E(G)| \geq (1 - \varepsilon)(2 + \delta)|V(G)| = (2 + \frac{\delta}{2})|V(G - F)|$. Therefore, $G - F$ has average degree $\geq (2 + \frac{\delta}{2})$ and girth $\geq 3 \cdot \frac{4k}{\delta}$, so by Lemma 2.17 it has a minor with average degree $\geq 2 + \frac{4k}{\delta} \cdot \frac{\delta}{2} > 2k$. But a minor of $G - F$ must have treewidth at most $\text{tw}(G - F) \leq k$, so average degree $\leq 2k$, a contradiction. \square

We now turn to classes of structures with a fixed signature σ . We will show that the class of tournaments (orientations of complete graphs) is not tw-pliable (or equivalently, size-pliable, by Theorem 1.5), in contrast to cliques and dense graphs (Example 7.2 and Theorem 1.3).

Lemma 2.19. *Let \mathcal{G} be a class of graphs of unbounded average degree. Let \mathcal{A} be the class of (unweighted) orientations of graphs in \mathcal{G} . Then \mathcal{A} is not size-pliable.*

In order to prove Lemma 2.19, we will need some definitions (only for this subsection) and alternative characterisations of size-pliability.

Definition 2.20. For a graph parameter p , let \bar{p} be the parameter defined as $\bar{p}(G) := \max_i(p(G_i))$, where G_i are the connected components of G .

For example, if p is size, then \bar{p} is cc (max component size). All the other parameters we consider (cc, td, tw, Hadwiger) satisfy $\bar{p} = p$.

Definition 2.21. A parameter p is *good* if \bar{p} -pliability is the same as size-pliability on classes of structures with bounded signatures.

We use this definition to state the next few lemmas in full generality. Theorem 1.5 shows that the parameters size, cc, td, tw, Hadwiger are good.

For two σ -structures \mathbb{A}, \mathbb{B} and a function $g: A \rightarrow B$, we define $\mathbb{I}m(g)$ to be the σ -structure on B with $f^{\mathbb{I}m(g)}(\mathbf{x}) := \min(f^{\mathbb{A}}(g^{-1}(\mathbf{x})), f^{\mathbb{B}}(\mathbf{x}))$. Note that $\mathbb{I}m(g) \subseteq \mathbb{B}$ (meaning each tuple has value in $\mathbb{I}m(g)$ less than or equal its value in \mathbb{B}).

Lemma 2.22. *Let p be a good parameter. Then a class of σ -structures \mathcal{A} is size-pliable if and only if $\forall \varepsilon > 0 \exists k \forall \mathbb{A} \in \mathcal{A}$ there is an overcast ω from \mathbb{A} to $(1 - \varepsilon)\mathbb{A}$ such that every $g: A \rightarrow A$ in its support has $p(\text{Im}(g)) \leq k$.*

Proof. For one direction, suppose that for every $\varepsilon > 0$ there is an integer k such that all $\mathbb{A} \in \mathcal{A}$ have an overcast ω from \mathbb{A} to $(1 - \varepsilon)\mathbb{A}$ such that every $g: A \rightarrow A$ in its support has $p(\text{Im}(g)) \leq k$. Then for these $\varepsilon, k, \mathbb{A}$ we can take \mathbb{B} to be the disjoint union of rescaled structures $\mathbb{B}_g := \omega(g) \text{Im}(g)$. We have $\bar{p}(\mathbb{B}) \leq k$. The overcast ω naturally induces overcasts showing $\mathbb{A} \succeq \mathbb{B} \succeq (1 - \varepsilon)\mathbb{A}$. Namely, we can define an overcast ω' from \mathbb{A} to \mathbb{B} by letting $\omega'(g') = \omega(g)$ for g' mapping A to $B_g \subseteq B$ just as g maps A to $\text{Im}(g) \subseteq \mathbb{A}$. We can also define an overcast ω'' from \mathbb{B} to $(1 - \varepsilon)\mathbb{A}$ by letting $\omega''(g'') = 1$ for one function g'' mapping each $\mathbb{B}_g \subseteq \mathbb{B}$ to $\text{Im}(g) \subseteq \mathbb{A}$. Hence $d_{\text{opt}}(\mathbb{A}, \mathbb{B}) \leq \varepsilon + \mathcal{O}(\varepsilon^2)$ (recall $1 \pm \varepsilon$ is close to $e^{\pm \varepsilon}$), which concludes the proof that \mathcal{A} is \bar{p} -pliable. Since we assume that p is a good parameter, \mathcal{A} is size-pliable.

In the other direction, suppose \mathcal{A} is size-pliable, meaning for every $\varepsilon > 0$ there is an integer k such that all \mathbb{A} have a \mathbb{B} with $d_{\text{opt}}(\mathbb{A}, \mathbb{B}) \leq \varepsilon$ and $|B| \leq k$. This means there are overcast ω and ω' showing $\mathbb{A} \succeq e^{-\varepsilon}\mathbb{B}$ and $\mathbb{B} \succeq e^{-\varepsilon}\mathbb{A}$, respectively. Then composing ω with ω' gives an overcast from \mathbb{A} to $(1 - 2\varepsilon)\mathbb{A}$ (since $e^{-2\varepsilon} \geq 1 - 2\varepsilon$), with the property that all images of functions g in the support are of size at most $|B| \leq k$, which implies $p(\text{Im}(g))$ is bounded by some function of k (namely $\max p(H)$ over all k -vertex graphs H). \square

We can now use Farkas' lemma to deduce another equivalent formulation:

Lemma 2.23. *Let p be a good parameter. Then a class of σ -structures \mathcal{A} is not size-pliable if and only if $\exists \varepsilon > 0 \forall k \in \mathbb{N}$ there is a pair of σ -structures $\mathbb{A} \in \mathcal{A}$ and \mathbb{C} with $C = A$, such that for every $g: A \rightarrow C$ with $p(\text{Im}(g)) \leq k$, $\text{value}(g) < (1 - \varepsilon) \text{value}(\text{id})$. (Here id is the identity map from A to $C = A$).*

Proof. By Lemma 2.22, \mathbb{A} is not size-pliable if and only if $\exists \varepsilon > 0 \forall k \in \mathbb{N}$ the following LP over variables $\{\omega(g): g \in V\}$, where $V := \{g \in A^A: p(\text{Im}(g)) \leq k\}$, has no non-negative rational solution:

$$\begin{aligned} \sum_{g \in V} \omega(g) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) &\geq (1 - \varepsilon) f^{\mathbb{A}}(\mathbf{x}) & \forall (f, \mathbf{x}) \in \text{tup}(\mathbb{A}) \\ \sum_{g \in V} \omega(g) &= 1 \end{aligned}$$

By applying Lemma A.2, this is equivalent to the existence of a non-negative vector $(y(f, \mathbf{x}))_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})}$ such that

$$\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} y(f, \mathbf{x}) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) < (1 - \varepsilon) \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A})} y(f, \mathbf{x}) f^{\mathbb{A}}(\mathbf{x}) \quad \forall g \in V$$

Let \mathbb{C} be the σ -structure on $C = A$ with $f^{\mathbb{C}}(\mathbf{x}) := y(f, \mathbf{x})$. Then the above inequality is restated as follows (interpreting $g \in V$ and id as maps from \mathbb{A} to \mathbb{C}):

$$\text{value}(g) < (1 - \varepsilon) \text{value}(\text{id}) \quad \forall g \in V. \quad \square$$

Remark 2.24. The structures \mathbb{A}, \mathbb{C} obtained above can be assumed to satisfy $G(\mathbb{A}) = G(\mathbb{C})$ without loss of generality, because for any $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$ such that one of $f^{\mathbb{A}}(\mathbf{x})$ or $f^{\mathbb{C}}(\mathbf{x})$ is zero, decreasing the other to zero will not change $\text{value}(\text{id})$ and can only decrease $\text{value}(g)$.

Proof of Lemma 2.19. Let ε be a constant to be chosen later ($\frac{1}{10}$ will do). Given any k , let $G \in \mathcal{G}$ be a graph with $m \geq 20 \cdot \binom{k}{2}$ edges, n vertices, and average degree $\frac{2m}{n} \geq 100 \log_2 k$. Let \mathbb{A} be a random orientation of G (each edge is independently oriented in either direction with probability $\frac{1}{2}$). We claim that with positive probability \mathbb{A} admits no map $g: A \rightarrow A$ to itself with image of size at most k such that $\text{value}(g) \geq (1 - \varepsilon) \text{value}(\text{id})$. This will prove that \mathcal{A} satisfies the conditions of Lemma 2.23 and hence is not size-pliable.

If a map as above existed, it would imply the existence of an oriented graph D (with at most one arc between every two vertices) on at most k vertices and a function $g: A \rightarrow V(D)$ with $\text{value}(g) \geq (1 - \varepsilon)m$. Observe that $\text{value}(g)$ is the number of arcs of \mathbb{A} that are correctly mapped by g (i.e., to an arc of D with the same orientation). Hence there would be a set F of at most εm arcs of \mathbb{A} such that g maps all arcs of $\mathbb{A} - F$ correctly. Let us bound the probability that there exists such D, F, g . The number of possible D is $\leq 3^{\binom{k}{2}}$; the number of possible F is $\leq \sum_{i=0}^{\varepsilon m} \binom{m}{i} \leq 2^{H(\varepsilon)m}$; the number of possible g is $\leq k^n$. Note that $\frac{2m}{n} \geq 100 \log_2 k$ and $3^{\binom{k}{2}} \leq 2^{m/10}$ by our choice of G . For fixed D, F, g , the probability that g maps all arcs of $\mathbb{A} - F$ correctly to D is at most $(\frac{1}{2})^{(1-\varepsilon)m}$. Hence in total the probability that some such D, F, g exist is at most

$$3^{\binom{k}{2}} \cdot 2^{H(\varepsilon)m} \cdot k^n \cdot 2^{-(1-\varepsilon)m} \leq 2^{n \log_2 k - (1-\varepsilon-H(\varepsilon)-\frac{1}{10})m} \leq 2^{-(1-\varepsilon-H(\varepsilon)-\frac{1}{10}-\frac{1}{50})m}.$$

This is less than 1 for ε small enough so that $1 - \varepsilon - H(\varepsilon) - \frac{1}{10} - \frac{1}{50} > 0$. \square

Finally, not all classes of bounded degree give pliable classes, even with a fixed signature.

Lemma 2.25. *Let \mathcal{G} be a class of graphs with unbounded girth and average degree $\geq 2 + \delta$ ($\delta > 0$). Let \mathcal{A} be the class of (unweighted) orientations of graphs in \mathcal{G} . Then \mathcal{A} is not size-pliable.*

Proof. We show there exists an ε such that for all k , there is an orientation $\mathbb{A} \in \mathcal{A}$ of a graph in \mathcal{G} such that every function $g: \mathbb{A} \rightarrow \mathbb{A}$ with $\text{cc}(\text{Im}(g)) \leq k$ has $\text{value}(g) < (1 - \varepsilon) \text{value}(\text{id})$. We choose ε later depending on δ only.

For any given k , let $G \in \mathcal{G}$ be a graph of girth $> k$. Let $m = |E(G)|$. Let \mathbb{A} be a random orientation of G : every edge is independently oriented in one direction or the other. We claim that the probability that there exists a $g: \mathbb{A} \rightarrow \mathbb{A}$ with $\text{cc}(\text{Im}(g)) \leq k$ and $\text{value}(g) \geq (1 - \varepsilon) \text{value}(\text{id})$ is strictly less than one (so there exists an orientation that satisfies our goal). Note that $\text{value}(\text{id}) = m$ and $\text{value}(g)$ is the number of arcs in \mathbb{A} that are mapped correctly (to an arc in \mathbb{A} with the same orientation); moreover, since the graph underlying \mathbb{A} has girth $> k$ and $\text{cc}(\text{Im}(g)) \leq k$, g must map into an oriented forest (disjoint union of trees). So the event is equivalent to the following: there exists a set $F \subseteq E(G)$ with $|F| \leq \varepsilon m$ and a function $g: \mathbb{A} \rightarrow \mathbb{A}$ which maps all arcs of $\mathbb{A} - F$ correctly into an oriented forest in \mathbb{A} .

The probability of this event can be union-bounded by the sum over $F \subseteq E(G)$ with $|F| \leq \varepsilon m$ of the probability that all of $\mathbb{A} - F$ can be mapped correctly into a subdigraph. The number of such F is $\sum_{i=0}^{\varepsilon m} \binom{m}{i} \leq 2^{H(\varepsilon)m}$; It remains to bound, for a fixed F , the probability that $\mathbb{A} - F$ can be mapped correctly.

Consider a fixed $F \subseteq E(G)$ with $|F| \leq \varepsilon m$. If $\mathbb{A} - F$ can be mapped correctly into an oriented forest in \mathbb{A} , then in particular it admits a homomorphism to C_3 , the directed cycle

digraph with three arcs. Let T be a spanning forest of $\mathbb{A} - F$ (a union of spanning trees of each connected component of $G - F$). There exists exactly one homomorphism from the edges of T in \mathbb{A} to C_3 (up to rotations in C_3 of each component); every remaining edge in $\mathbb{A} - F - E(T)$ closes an oriented cycle, so it has at most one orientation which allows to extend this unique homomorphism to it. Hence the probability that $\mathbb{A} - F$ admits a homomorphism to C_3 is at most $(\frac{1}{2})^{m'}$ where $m' = m - |F| - |E(T)| \geq m - \varepsilon m - |V(G)| \geq (1 - \varepsilon - \frac{2}{2+\delta})m = (\frac{\delta}{2+\delta} - \varepsilon)m$.

All in all, the probability of our original event is at most $2^{H(\varepsilon) \cdot m} \cdot (\frac{1}{2})^{m'} \leq 2^{-(\frac{\delta}{2+\delta} - \varepsilon - H(\varepsilon))m}$. Hence it suffices to choose ε small enough so that $\varepsilon + H(\varepsilon) < \frac{\delta}{2+\delta}$. \square

3 Pliable structures admit a PTAS: proof of Theorem 1.1

We first define the Sherali-Adams LP hierarchy [83] for Max-Hom. Let (\mathbb{A}, \mathbb{B}) be an instance of Max-Hom over a signature σ and let $k \geq \max_{f \in \sigma} \text{ar}(f)$. For a tuple \mathbf{x} , we denote by $\text{Set}(\mathbf{x})$ the set of elements appearing in \mathbf{x} . 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* [83] of (\mathbb{A}, \mathbb{B}) is the linear program given in Figure 1, which has one variable $\lambda(X, s)$ for each $X \in \binom{A}{\leq k}$ and each $s: X \rightarrow B$.

We denote by $\text{opt}_k(\mathbb{A}, \mathbb{B})$ the optimum value of this linear program.

$$\begin{array}{ll}
\max & \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{A}), s: \text{Set}(\mathbf{x}) \rightarrow B} \lambda(\text{Set}(\mathbf{x}), s) f^{\mathbb{A}}(\mathbf{x}) f^{\mathbb{B}}(s(\mathbf{x})) \\
& \lambda(X, s) = \sum_{r: Y \rightarrow B, r|_X = s} \lambda(Y, r) \quad \text{for } X \subseteq Y \in \binom{A}{\leq k} \text{ and } s: X \rightarrow B \\
& \sum_{s: X \rightarrow B} \lambda(X, s) = 1 \quad \text{for } X \in \binom{A}{\leq k} \\
& \lambda(X, s) \geq 0 \quad \text{for } X \in \binom{A}{\leq k} \text{ and } s: X \rightarrow B
\end{array}$$

Figure 1: The Sherali-Adams relaxation of level k of Max-Hom instance (\mathbb{A}, \mathbb{B}) .

Observation 3.1. Let \mathbb{A} be a σ -structure, $\lambda \geq 0$ and $k \geq \max_{f \in \sigma} \text{ar}(f)$. Then for all σ -structures \mathbb{C} , we have $\text{opt}(\lambda \mathbb{A}, \mathbb{C}) = \lambda \text{opt}(\mathbb{A}, \mathbb{C})$ and $\text{opt}_k(\lambda \mathbb{A}, \mathbb{C}) = \lambda \text{opt}_k(\mathbb{A}, \mathbb{C})$.

Definition 3.2. Let \mathbb{A} and \mathbb{B} be σ -structures and $k \geq \max_{f \in \sigma} \text{ar}(f)$. We write $\mathbb{A} \succeq_k \mathbb{B}$ if, for all σ -structures \mathbb{C} , we have $\text{opt}_k(\mathbb{A}, \mathbb{C}) \geq \text{opt}_k(\mathbb{B}, \mathbb{C})$.

The proof of the following is analogous to the proof of [11, Proposition 5.2]. For completeness, it is given in Appendix B.

Proposition 3.3. Let \mathbb{A} and \mathbb{B} be σ -structures and $k \geq \max_{f \in \sigma} \text{ar}(f)$. If there is an overcast from \mathbb{A} to \mathbb{B} then $\mathbb{A} \succeq_k \mathbb{B}$.

Using Observation 3.1 and Proposition 3.3, we are ready to prove the following.

Proposition 3.4. Let \mathbb{A} be a σ -structure. Let $\varepsilon \geq 0$ and $k \geq \max_{f \in \sigma} \text{ar}(f)$. Suppose that there exists a 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 σ -structure \mathbb{C} , we have that

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

Proof. The left-hand side inequality is from the definition of Sherali-Adams. For the right-hand side inequality, observe first that, by definition of d_{opt} , $\mathbb{A} \succeq e^{-\varepsilon} \mathbb{B}$ and $\mathbb{B} \succeq e^{-\varepsilon} \mathbb{A}$. By Proposition 2.5, there is an overcast from \mathbb{B} to $e^{-\varepsilon} \mathbb{A}$, so by Proposition 3.3, it follows that $\mathbb{B} \succeq_k e^{-\varepsilon} \mathbb{A}$. By Observation 3.1, we have that $\text{opt}_k(\mathbb{B}, \mathbb{C}) \geq e^{-\varepsilon} \text{opt}_k(\mathbb{A}, \mathbb{C})$. Since $\text{tw}(\mathbb{B}) \leq k$, we have $\text{opt}_k(\mathbb{B}, \mathbb{C}) = \text{opt}(\mathbb{B}, \mathbb{C})$ – this follows, for example, from [11, Theorem 5.8].⁸ Since moreover $\mathbb{A} \succeq e^{-\varepsilon} \mathbb{B}$, by Observation 3.1, it follows that $\text{opt}(\mathbb{A}, \mathbb{C}) \geq e^{-\varepsilon} \text{opt}(\mathbb{B}, \mathbb{C})$. Together, $\text{opt}(\mathbb{A}, \mathbb{C}) \geq e^{-\varepsilon} \text{opt}(\mathbb{B}, \mathbb{C}) = e^{-\varepsilon} \text{opt}_k(\mathbb{B}, \mathbb{C}) \geq e^{-2\varepsilon} \text{opt}_k(\mathbb{A}, \mathbb{C})$. Hence $\text{opt}_k(\mathbb{A}, \mathbb{C}) \leq e^{2\varepsilon} \text{opt}(\mathbb{A}, \mathbb{C})$. \square

Since $\text{opt}_k(\mathbb{A}, \mathbb{C})$ can be computed in time $(|A| \cdot |C|)^{\mathcal{O}(k)}$, this concludes the proof of Theorem 1.1.

4 Fractional fragility

To give Dvořák’s definition of fractional fragility [29] we first define ε -thin distributions.

Definition 4.1. Let \mathcal{F} be a family of subsets of a set V and $\varepsilon > 0$. We say that 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 4.2. 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* if for every $\varepsilon > 0$ there is a k such that every $G \in \mathcal{G}$ has an ε -thin fractional $(p \leq k)$ -modulator. We can analogously define $(p \leq k)$ -edge-modulators $F \subseteq E(G)$ and *fractionally- p -edge-fragility*.

One crucial property of fractional fragility is that it allows a dual definition by a variant of Farkas’ lemma (cf. Appendix A for details); this is already implicit in [32, Lemma 6].

Lemma 4.3. Let \mathcal{F} be a family of subsets of a set V . The following are equivalent:

- there is an ε -thin distribution π of sets in \mathcal{F} ;
- for all non-negative weights $(w(v))_{v \in V}$, there is an $X \in \mathcal{F}$ such that $w(X) := \sum_{x \in X} w(x) \leq \varepsilon \cdot w(V)$.

Thus a class of graphs \mathcal{G} is fractionally-tw-fragile if and only if for every $\varepsilon > 0$ there is a k such that for every graph $G \in \mathcal{G}$ and every vertex-weight function w , one can remove a set of vertices of weight at most $\varepsilon \cdot w(V)$ to obtain a graph with $\text{tw} \leq k$.

Another useful property of fractional fragility is that the edge version is equivalent to the vertex version, for most parameters of interest. Recall that each parameter we consider (Hadwiger, tw , tw , td , cc , size) is *monotone*, meaning $p(H) \leq p(G)$ for H a subgraph of G ; and that the average degree $\frac{2|E(G)|}{|V(G)|}$ is bounded by a function of p .

Lemma 4.4. Let p be a monotone graph parameter such that the average degree $\frac{2|E(G)|}{|V(G)|}$ of a graph is bounded by a function of $p(G)$. Let \mathcal{G} be a class of graphs. Then the following are equivalent:

⁸Our definition of the LP slightly differs from [11], where there are additional variables $\lambda(f, \mathbf{x}, s)$ associated with tuples (f, \mathbf{x}) with $f^{\mathbf{A}}(\mathbf{x}) > 0$. However, since we are assuming without loss of generality that $k \geq \max_{f \in \sigma} \text{ar}(f)$, the two definitions are equivalent.

- \mathcal{G} is fractionally-p-fragile;
- \mathcal{G} is fractionally-p-edge-fragile;
- $\forall \varepsilon > 0 \exists k \forall G \in \mathcal{G} \forall w: V(G) \rightarrow \mathbb{Q}_{\geq 0} \exists X \subseteq V(G) \ w(X) \leq \varepsilon w(V(G))$ and $p(G - X) \leq k$;
- $\forall \varepsilon > 0 \exists k \forall G \in \mathcal{G} \forall w: E(G) \rightarrow \mathbb{Q}_{\geq 0} \exists F \subseteq E(G) \ w(F) \leq \varepsilon w(E(G))$ and $p(G - F) \leq k$.

Proof. (i) is equivalent to (iii) and (ii) is equivalent to (iv) by Lemma 4.3.

It is easy to see that (i) implies (iv): suppose for every $\varepsilon > 0$ there is a k such that every $G \in \mathcal{G}$ has an ε -thin fractional $(p \leq k)$ -modulator π . Let $w: E(G) \rightarrow \mathbb{Q}_{\geq 0}$ be any edge-weight function. If we take a set X from the distribution π and remove the set F of all edges incident to X , this yields a graph with $p(G - F) \leq k$. Every vertex is in X with probability $\leq \varepsilon$, so every edge is in F with probability $\leq 2\varepsilon$. Hence the expected weight of F is $\leq 2\varepsilon w(E(G))$. So there exists a set $F \subseteq E(G)$ such that $w(F) \leq 2\varepsilon w(E(G))$ and $p(G - F) \leq k$.

It remains to show that (iv) implies (iii). Let $f: \mathbb{N} \rightarrow \mathbb{N}$ be such that $\frac{2|E(G)|}{|V(G)|} \leq f(p(G))$ for all graphs G .

We first show that (iv) implies that \mathcal{G} has bounded *maximum average degree* $\text{mad}(G) := \max_{H \subseteq G} \frac{2|E(H)|}{|V(H)|}$. Indeed, let $k := k(\varepsilon)$ be a number satisfying (iv) for $\varepsilon = \frac{1}{2}$. Then for any $G \in \mathcal{G}$ and any $H \subseteq G$, let $w: E(G) \rightarrow \mathbb{Q}_{\geq 0}$ assign 1 to edges in H and 0 to edges not in H . By assumption there is a set $F \subseteq E(G)$ such that $w(F) \leq \varepsilon w(E(G))$ and $p(G - F) \leq k$. Let $F' := F \cap E(H)$; then $|F'| = w(F) \leq \varepsilon w(E(G)) = \varepsilon |E(H)|$ and $p(H - F') \leq p(G - F) \leq k$. Hence $(1 - \varepsilon)|E(H)| \leq |E(H - F')| \leq \frac{f(k)}{2} \cdot |V(H - F')|$, which means $\frac{2|E(H)|}{|V(H)|} \leq \frac{f(k(\frac{1}{2}))}{1 - \varepsilon} = 2f(k(\frac{1}{2}))$. That is, every subgraph H of every graph G in \mathcal{G} has average degree at most $D := 2f(k(\frac{1}{2}))$.

This implies that every subgraph has some vertex of degree at most D (this is called the *degeneracy* of the graph: it is upper bounded by mad). Hence every graph G in \mathcal{G} has an orientation \vec{G} with maximum in-degree at most D (obtained by iteratively finding a vertex of degree at most D , orienting all remaining edges towards it, and removing the vertex).

To show (iii), let $\varepsilon > 0$, $k' := k(\frac{\varepsilon}{D})$, $G \in \mathcal{G}$. Choose an orientation \vec{G} of G with maximum in-degree at most D . Given $w: V(G) \rightarrow \mathbb{Q}_{\geq 0}$, we can define $w': E(G) \rightarrow \mathbb{Q}_{\geq 0}$ as $w'(uv) := w(v)$ if uv is directed towards v . By assumption, there is a set of edges F such that $p(G - F) \leq k'$ and

$$w'(F) \leq \frac{\varepsilon}{D} w'(E(G)).$$

Let $X := \{v: \exists uv \in F \text{ directed towards } v\}$; then $G - X \subseteq G - F$, so $p(G - X) \leq k'$. Note that

$$w'(E(G)) = \sum_{\vec{uv} \in E(\vec{G})} w(v) = \sum_{v \in V(G)} \text{in-deg}(v) \cdot w(v) \leq D \cdot w(V(G))$$

and

$$w'(F) = \sum_{\vec{uv} \in F} w(v) \geq \sum_{v \in X} w(v) = w(X)$$

Hence

$$w(X) \leq w'(F) \leq \frac{\varepsilon}{D} w'(E(G)) \leq \varepsilon \cdot w(V(G)).$$

This concludes the proof that (iv) implies (iii). \square

Dvořák and Sereni [32, Theorem 31] showed that graphs of bounded treewidth are fractionally-td-fragile. It follows from a result of DeVos et al. [24, Theorem 1.2] that for H -minor-free graphs (for any graph H , that is classes of graphs that exclude some minor, or equivalently, classes of bounded Hadwiger number) are fractionally-tw-fragile.⁹ These two facts establish the following equivalence, by [29, Lemma 12].

Theorem 4.5 ([32, 24, 29]). *The following are equivalent for a class of graphs \mathcal{G} :*

- \mathcal{G} is fractionally-td-fragile;
- \mathcal{G} is fractionally-tw-fragile;
- \mathcal{G} is fractionally-Hadwiger-fragile.

4.1 Fragility implies pliability: proof of Theorem 1.2

We denote by $G \uplus H$ the disjoint union of graphs G and H . All graph parameters p we consider satisfy $p(G \uplus H) = \max(p(G), p(H))$ for all G, H (that is: cc, td, tw, Hadwiger, excluding only size; we never consider fractional-size-fragility, as it is equivalent to just bounded size).

Lemma 4.6. *Let p be a graph parameter such that $p(G \uplus H) = \max(p(G), p(H))$ for all G, H . Let \mathcal{A} be a class of structures of bounded arity r such that the class \mathcal{G} of their Gaifman graphs is fractionally- p -fragile. Then \mathcal{A} is p -pliable.*

Proof. For $\varepsilon > 0$, let $\varepsilon' := \frac{\varepsilon}{1+\varepsilon} \cdot \frac{1}{r}$. By definition of fractional- p -fragility, $\exists_{k(\varepsilon)} \forall_{G \in \mathcal{G}} G$ has an ε' -thin fractional ($p \leq k$)-modulator. Let $\mathbb{A} \in \mathcal{A}$ be a structure with Gaifman graph $G \in \mathcal{G}$. By assumption, G has a fractional ($p \leq k$)-modulator π such that for every $v \in V(G)$, $\Pr_{X \sim \pi}[v \in X] \leq \varepsilon'$. For $X \subseteq V(G) = A$ in the support of π , let \mathbb{B}_X be the rescaling of $\mathbb{A} - X$ by a factor of $\pi(X)$; let \mathbb{B} be the disjoint union of all \mathbb{B}_X . Since each X in the support of π is a ($p \leq k$)-modulator and p is closed under disjoint union, $p(G(\mathbb{B})) \leq k$.

We define overcasts $\omega: \mathbb{A} \rightarrow \mathbb{B}$ and $\omega': \mathbb{B} \rightarrow (1 - r\varepsilon')\mathbb{A}$. The first, ω , maps \mathbb{A} identically to each component \mathbb{B}_X of \mathbb{B} with probability $\pi(X)$ (vertices of \mathbb{A} in X are mapped arbitrarily in the same component). The second, ω' , deterministically maps each component \mathbb{B}_X of \mathbb{B} identically to \mathbb{A} . To check that ω' is indeed an overcast, consider a tuple $(f, \mathbf{x}) \in \text{tup}(\mathbb{A})$. The tuple is covered by its copies in \mathbb{B}_X with weight $\pi(X) \cdot f^{\mathbb{A}}(X)$ for all X which do not intersect \mathbf{x} . In total, the fraction of $f^{\mathbb{A}}(\mathbf{x})$ lost is hence exactly $\Pr_{X \sim \pi}[X \cap \mathbf{x} \neq \emptyset]$, which is (by union bound and by the assumption $|\mathbf{x}| \leq r$) at most $\varepsilon'r$. Since $1 - \varepsilon'r = \frac{1}{1+\varepsilon} \geq e^{-\varepsilon}$, we have $\mathbb{A} \succeq \mathbb{B} \succeq (1 - \varepsilon'r)\mathbb{A} \succeq e^{-\varepsilon}\mathbb{A}$, which means \mathbb{B} is a structure at opt-distance $\leq \varepsilon$ from \mathbb{A} . \square

This concludes Theorem 1.2: structures on fractionally-tw-fragile graphs are tw-pliable.

4.2 Pliability vs fragility: proof of Lemma 1.4

For Lemma 1.4, we need the other direction than in Theorem 1.2: that if all structures on Gaifman graphs in \mathcal{G} are tw-pliable, then \mathcal{G} is fractionally-tw-fragile. To do this, intuitively, we consider, for a graph $G \in \mathcal{G}$, a structure \mathbb{A} where each edge is used by a different symbol of the signature. If we have a structure \mathbb{B} (of bounded treewidth) close to \mathbb{A} in opt-distance, this

⁹In fact, as shown by Dvořák [30], a proof of van den Heuvel et al. [88, Lemma 4.1] can be adapted to show this without the Graph Minors Structure Theorem.

implies overcasts from \mathbb{A} to $e^{-\varepsilon}\mathbb{B}$ and from \mathbb{B} to $e^{-\varepsilon}\mathbb{A}$; composing the two gives an overcast from $e^{+\varepsilon}\mathbb{A}$ to $e^{-\varepsilon}\mathbb{A}$ in which (since each edge is used by a different symbol) an edge can only be covered by itself. This shows that the overcasts are mostly injective and that \mathbb{B} , sandwiched between $e^{+\varepsilon}\mathbb{A}$ and $e^{-\varepsilon}\mathbb{A}$, must be close in edit distance. The bounded treewidth of \mathbb{B} then implies that the graph G underlying \mathbb{A} is in fact fractionally-tw-edge-fragile, which by Lemma 4.4 concludes the proof.

The formal proof of Lemma 1.4 follows. In fact, we prove the statement for any reasonable parameter, including cc, td, tw, and Hadwiger; the conclusion is the edge variant of fractional fragility, but the two are equivalent by Lemma 4.4.

Lemma 4.7 (Lemma 1.4 more generally). *Let p be a monotone graph parameter such that $p(G \uplus H) = \max(p(G), p(H))$. For every integer $r \geq 2$, a class of graphs \mathcal{G} is fractionally-p-edge-fragile if and only if $\mathcal{A}_{\mathcal{G}}^{(r)}$ is p -pliable.*

Proof. If \mathcal{G} is fractionally-p-fragile, then by Lemma 4.6 $\mathcal{A}_{\mathcal{G}}^{(r)}$ is p -pliable.

For the other direction, suppose $\mathcal{A}_{\mathcal{G}}^{(r)}$ is p -pliable:

$$\forall \varepsilon > 0 \exists k(\varepsilon) \forall \mathbb{A} \in \mathcal{A}_{\mathcal{G}}^{(r)} \exists \mathbb{B} \quad p(\mathbb{B}) \leq k \text{ and } d_{\text{opt}}(\mathbb{A}, \mathbb{B}) \leq \varepsilon.$$

Let $\varepsilon > 0$ and let $k := k(\frac{\varepsilon}{2})$. For a graph $G \in \mathcal{G}$, let σ be the signature with a different binary symbol f_e for each $e \in E(G)$. Let \mathbb{A} be the σ -structure with domain $V(G)$ and values $f_e^{\mathbb{A}}(u, v) = 1$ if $\{u, v\} = e$, 0 otherwise. (The arity can be increased to exactly r by adding dummy or repeated variables). By assumption, there is a σ -structure \mathbb{B} such that $p(\mathbb{B}) \leq k$ and $d_{\text{opt}}(\mathbb{A}, \mathbb{B}) \leq \frac{\varepsilon}{2}$. Let ω, ω' be overcasts from \mathbb{A} to $\exp(-\frac{\varepsilon}{2}) \cdot \mathbb{B}$ and from \mathbb{B} to $\exp(-\frac{\varepsilon}{2}) \cdot \mathbb{A}$, respectively.

For g with $\omega(g) > 0$ and g' with $\omega'(g') > 0$, let $F_{gg'} \subseteq E(G)$ be the subset of edges e such that $g'(g(e)) \neq e$ or $f_e^{\mathbb{B}}(g(e)) = 0$. Since $g' \circ g$ is the identity on $E(G) - F_{gg'}$, the functions g' and g are bijections between this set and a subset of edges of $G(\mathbb{B})$. Hence $G - F_{gg'}$ is isomorphic to a subgraph of $G(\mathbb{B})$, which implies $p(G - F_{gg'}) \leq k$.

Let $e \in E(G)$. We claim that $\Pr_{\substack{g \sim \omega \\ g' \sim \omega'}} [e \in F_{gg'}] \leq \varepsilon$. This holds essentially because the composition of ω and ω' is an overcast from \mathbb{A} to $\exp(-\varepsilon) \cdot \mathbb{A}$ and because the only edge with non-zero value of $f_e^{\mathbb{A}}$ is e itself. Formally, since ω is an overcast, we have:

$$\text{for each } e_B \in E(G(\mathbb{B})) \quad \mathbb{E}_{g \sim \omega} f_e^{\mathbb{A}}(g^{-1}(e_B)) \geq \exp(-\frac{\varepsilon}{2}) f_e^{\mathbb{B}}(e_B).$$

Note that by construction of \mathbb{A} , $f_e^{\mathbb{A}}(g^{-1}(e_B)) = [g(e) = e_B]$. Hence for each $e_B \in E(G(\mathbb{B}))$,

$$\mathbb{E}_{g \sim \omega} [g(e) = e_B] \geq \exp(-\frac{\varepsilon}{2}) f_e^{\mathbb{B}}(e_B).$$

Moreover, since ω' is an overcast, we have:

$$\mathbb{E}_{g' \sim \omega'} f_e^{\mathbb{B}}(g'^{-1}(e)) \geq \exp(-\frac{\varepsilon}{2}) f_e^{\mathbb{A}}(e).$$

That is:

$$\mathbb{E}_{g' \sim \omega'} \sum_{\substack{e_B \in E(G(\mathbb{B})) \\ g'(e_B) = e}} f_e^{\mathbb{B}}(e_B) \geq \exp(-\frac{\varepsilon}{2}).$$

Putting the two together:

$$\begin{aligned}
\Pr_{\substack{g \sim \omega \\ g' \sim \omega'}} [e \notin F_{gg'}] &= \Pr_{\substack{g \sim \omega \\ g' \sim \omega'}} [g'(g(e)) = e \text{ and } g(e) \in E(G(\mathbb{B}))] \\
&= \mathbb{E}_{\substack{g \sim \omega \\ g' \sim \omega'}} \sum_{\substack{e_B \in E(G(\mathbb{B})) \\ g'(e_B) = e}} [g(e) = e_B] \\
&\geq \mathbb{E}_{\substack{g \sim \omega \\ g' \sim \omega'}} \sum_{\substack{e_B \in E(G(\mathbb{B})) \\ g'(e_B) = e}} \exp(-\frac{\varepsilon}{2}) f_e^{\mathbb{B}}(e_B) \\
&\geq \exp(-\varepsilon) \geq 1 - \varepsilon.
\end{aligned}$$

Therefore, we obtained a distribution of edge sets $F_{gg'} \subseteq E(G)$ such that $p(G - F_{gg'}) \leq k$ satisfying $\Pr_{\substack{g \sim \omega \\ g' \sim \omega'}} [e \in F_{gg'}] \leq \varepsilon$. This is an ε -thin fractional $(p \leq k)$ -edge-modulator. \square

5 From Hadwiger- to size-pliability: proof of Theorem 1.5

In this section, we prove Theorem 1.5: pliability with respect to different parameters yields equivalent definitions. The first half of Theorem 1.5 follows easily from already established results and a simple observation, cf. Section 5.1. The second half of Theorem 1.5 reduces to showing that structures of bounded treedepth with a bounded signature are size-pliable. The strategy for the proof is similar to a proof of Nešetřil and Ossona de Mendez [74, Corollary 3.3] that relational structures of bounded treedepth have bounded cores. However the argument is much more intricate due to the fact that we consider valued structures: the statement that there are only finitely many structures of size at most C , for every C , is not true anymore. The main difficulty is proving an approximate version of it.

5.1 Treewidth-, treedepth-, and Hadwiger-pliability

The first half of Theorem 1.5, that is, the equivalence of p -pliability for $p \in \{\text{tw}, \text{td}, \text{Hadwiger}\}$, will follow (as detailed in Corollary 5.2 below) from the equivalence of fractional- p -fragility for these parameters (Theorem 4.5), the fact that fragility implies pliability (Lemma 4.6), and transitivity of pliability, in the following sense.

Observation 5.1 (Transitivity of pliability). *Let \mathcal{A} be a class of structures with signatures from a set Σ . Suppose \mathcal{A} is p -pliable and for each k , $\{\mathbb{A}: p(\mathbb{A}) \leq k\}$ is p' -pliable, where \mathbb{A} runs over all structures with signatures in Σ . Then \mathcal{A} is p' -pliable.*

Proof. Intuitively, this holds because d_{opt} is a pseudometric. Formally, suppose a class \mathcal{A} is p -pliable. Then every $\mathbb{A} \in \mathcal{A}$ is $\frac{\varepsilon}{2}$ -close (in d_{opt} distance) to some \mathbb{B} with $p(\mathbb{B}) \leq k$ (for some k depending on $\frac{\varepsilon}{2}$). By assumption, every \mathbb{B} with $p(\mathbb{B}) \leq k$ is $\frac{\varepsilon}{2}$ -close to some \mathbb{C} with $p'(\mathbb{C}) \leq k'$ (for some k' depending on $\frac{\varepsilon}{2}$ and k). Hence \mathbb{A} is ε -close to some structure \mathbb{C} with $p'(\mathbb{C}) \leq k'(\frac{\varepsilon}{2}, k(\frac{\varepsilon}{2}))$, which only depends on ε . \square

Corollary 5.2. *Let \mathcal{A} be any class of structures. The following are equivalent:*

- \mathcal{A} is td -pliable;

- \mathcal{A} is *tw-pliable*;
- \mathcal{A} is *Hadwiger-pliable*.

Proof. Since $\text{td}(G) \geq \text{tw}(G) + 1 \geq \text{Hadwiger}(G)$ for any graph G , each bullet point implies the next by Observation 2.9. It suffices to show that Hadwiger-pliability implies td-pliability. By Observation 5.1 it suffices to show that for every k , the class \mathcal{A} of all structures with Hadwiger number at most k (and arbitrary signatures) is td-pliable. These are structures whose Gaifman graphs exclude the clique K_{k+1} as a minor. Their Gaifman graphs are thus fractionally-td-fragile by Theorem 4.5. Since their Gaifman graphs do not include cliques K_{k+1} the arity of symbols with non-zero tuples is bounded by k . By Lemma 4.6, this implies that \mathcal{A} is td-pliable (high-arity symbols with no non-zero tuples can be ignored). \square

5.2 From cc-pliability to size-pliability

To show the second half of Theorem 1.5, i.e., the equivalence of td-pliability, cc-pliability, and size-pliability (for structures with bounded signatures), it will be easier to first focus on the latter two. Since there are only finitely many distinct signatures of bounded size and arity, we can focus on a single fixed signature (as finite unions of pliable classes are pliable).

Since $\text{cc} \leq \text{size}$, by Observation 2.9 we have that size-pliability implies cc-pliability. The rest of this section is devoted to proving that cc-pliability implies size-pliability (for a single fixed signature). This (and in fact equivalence of the two) would be trivial if there were only a bounded number of distinct values of tuples, since then there can be only a bounded number of components up to isomorphism, and isomorphic components can be merged.

Observation 5.3. *For any structure \mathbb{A} and numbers $\lambda_1, \dots, \lambda_n \in \mathbb{Q}_{\geq 0}$, the disjoint union $\lambda_1 \mathbb{A} \uplus \dots \uplus \lambda_n \mathbb{A}$ is equivalent (i.e., at d_{opt} -distance zero) to $\lambda \mathbb{A}$, where $\lambda = \lambda_1 + \dots + \lambda_n$.*

Proof. An overcast in one direction deterministically maps each component $\lambda_i \mathbb{A}$ to $\lambda \mathbb{A}$ identically. An overcast in the other direction maps $\lambda \mathbb{A}$ to the component $\lambda_i \mathbb{A}$ with probability λ_i / λ . \square

For a structure \mathbb{A} with components of bounded size and $\mathbb{Q}_{\geq 0}$ -values, we can try to change the values slightly to find a structure \mathbb{B} at small edit distance which uses a bounded number of different values (and then proceed as above). This works if the ratio of the maximum value to the minimum non-zero value is bounded. If this ratio is large, we could try to change the extremely small values to zero, hoping the edit distance is small (relative to the extremely large values). However, this does not always work: consider structures \mathbb{A} with few large values and many small values (for example a structure having 2^i tuples of value 2^{n-i} , for $i = 0 \dots n$). So the general case cannot be reduced to the case of finitely many distinct values just by finding a structure close in edit distance. Nevertheless, instead of requiring the modified structure \mathbb{B} to have a bounded number of components up to isomorphism, it suffices to require a bounded number of components up to rescaling (two structures $\mathbb{B}_1, \mathbb{B}_2$ being the same up to rescaling if $\mathbb{B}_1 = \lambda \mathbb{B}_2$ for some $\lambda > 0$). This minor weakening turns out to be sufficient to fix our problem. We formalise this first as a statement on sequences of vectors of bounded dimension (which will encode a sequence of components of bounded size).

Lemma 5.4. *Let $d \in \mathbb{N}, \varepsilon > 0$. There is a k such that for every sequence of vectors $v^{(1)}, \dots, v^{(n)} \in \mathbb{Q}_{\geq 0}^d$, there is a sequence $w^{(1)}, \dots, w^{(n)} \in \mathbb{Q}_{\geq 0}^d$ such that for each coordinate*

$i = 1, \dots, d$, $\sum_{j=1\dots n} |v_i^{(j)} - w_i^{(j)}| \leq \varepsilon \sum_{j=1\dots n} v_i^{(j)}$, and such that up to rescaling, there are only k distinct vectors in $w^{(1)}, \dots, w^{(n)}$.

Proof. The proof is by induction on d . Let $d \in \mathbb{N}, \varepsilon > 0$ and consider a sequence $v_1, \dots, v_n \in \mathbb{Q}_{\geq 0}^d$. For $d = 1$, the sequence already has only one vector up to rescaling (or two, if it contains the zero vector), so let $d \geq 2$.

Let $J = \{1, \dots, n\}$. For a subset $X \subseteq J$, denote $\text{mass}_i(X) := \sum_{j \in X} v_i^{(j)}$. We focus on the first two coordinates and in particular the ratio of the second to the first. For $c \in \mathbb{R}$, let $J_{<c} := \{j \in J : v_2^{(j)} < c \cdot v_1^{(j)}\}$. Define $J_{\leq c}, J_{>c}, J_{\geq c}$ analogously.

Let c be maximum such that $\text{mass}_2(J_{<c}) \leq \frac{\varepsilon}{3} \cdot \text{mass}_2(J)$. For $j \in J_{<c}$, let w_j be the vector obtained from v_j by zeroing the 2nd coordinate. The resulting difference is

$$\sum_{j \in J_{<c}} |v_2^{(j)} - w_2^{(j)}| \leq \frac{\varepsilon}{3} \sum_{j \in J} v_2^{(j)}.$$

By maximality of c we have

$$\text{mass}_2(J_{\leq c}) > \frac{\varepsilon}{3} \cdot \text{mass}_2(J).$$

Observe that the left hand side can be bounded as follows:

$$\text{mass}_2(J_{\leq c}) \leq c \cdot \text{mass}_1(J_{\leq c}) \leq c \cdot \text{mass}_1(J)$$

and similarly the right hand side can be bounded as follows, for $c' := c \cdot \frac{3d}{\varepsilon^2}$:

$$\frac{\varepsilon}{3} \cdot \text{mass}_2(J) \geq \frac{\varepsilon}{3} \cdot \text{mass}_2(J_{\geq c'}) \geq c \cdot \frac{d}{\varepsilon} \cdot \text{mass}_1(J_{\geq c'}).$$

Altogether, this implies

$$c \cdot \text{mass}_1(J) > c \cdot \frac{d}{\varepsilon} \cdot \text{mass}_1(J_{\geq c'}),$$

which after rearranging gives

$$\text{mass}_1(J_{\geq c'}) < \frac{\varepsilon}{d} \cdot \text{mass}_1(J).$$

For $j \in J_{\geq c'}$, let w_j be the vector obtained from v_j by zeroing the 1st coordinate. The resulting difference is $\sum_{j \in J_{\geq c'}} |v_1^{(j)} - w_1^{(j)}| \leq \frac{\varepsilon}{d} \sum_{j \in J} v_1^{(j)}$.

The only remaining vectors, in $J_{\text{mid}} := J \setminus (J_{<c} \cup J_{\geq c'})$, satisfy $c \cdot v_1^{(j)} \leq v_2^{(j)} < c' \cdot v_1^{(j)}$. We can round down their 2nd coordinate to $c \cdot v_1^{(j)}$ times an integer power of $e^{\varepsilon/3}$. That is, for $j \in J_{\text{mid}}$, let w_j be the vector obtained from v_j by decreasing the 2nd coordinate to $w_2^{(j)} := c \cdot e^{a\varepsilon/3} \cdot v_1^{(j)}$ with $a \in \mathbb{N}$ maximum such that $w_2^{(j)} \leq v_2^{(j)}$. Observe that $a \geq 0$ and since $c \cdot e^{a\varepsilon/3} \cdot v_1^{(j)} \leq v_2^{(j)} \leq c' \cdot v_1^{(j)}$, we have $e^{a\varepsilon/3} \leq \frac{c'}{c} = \frac{3d}{\varepsilon^2}$ and thus $a \leq \frac{3}{\varepsilon} \cdot \ln(\frac{3d}{\varepsilon^2})$. Note also that $1 \geq \frac{w_2^{(j)}}{v_2^{(j)}} > e^{-\varepsilon/3}$, hence $\frac{|v_2^{(j)} - w_2^{(j)}|}{v_2^{(j)}} \leq 1 - e^{-\varepsilon/3} < \frac{\varepsilon}{3}$, so the resulting difference is

$$\sum_{j \in J_{\text{mid}}} |v_2^{(j)} - w_2^{(j)}| \leq \frac{\varepsilon}{3} \sum_{j \in J} v_2^{(j)}$$

To summarise, all vectors w_j satisfy $w_j \leq v_j$ (coordinate-wise) and when limited to their first two coordinates as $\begin{pmatrix} w_1^{(j)} \\ w_2^{(j)} \end{pmatrix}$, are either multiples of $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ (if $j \in J_{<c}$), or multiples of $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$

(if $j \in J_{\geq c'}$), or multiples of $(\frac{1}{c \cdot e^{a\varepsilon/3}})$, for some $a \in \{0, 1, \dots, K\}$ for $K := \lfloor \frac{3}{\varepsilon} \cdot \ln(\frac{3d}{\varepsilon^2}) \rfloor$. The resulting differences in the first and second coordinate, respectively, are bounded by as

$$\sum_{j \in J_{\geq c'}} |v_1^{(j)} - w_1^{(j)}| \leq \frac{\varepsilon}{d} \sum_{j \in J} v_1^{(j)}$$

$$\sum_{j \in J_{< c}} |v_2^{(j)} - w_2^{(j)}| + \sum_{j \in J_{mid}} |v_2^{(j)} - w_2^{(j)}| \leq (\frac{\varepsilon}{3} + \frac{\varepsilon}{3}) \sum_{j \in J} v_2^{(j)}.$$

We replace the sequence $v^{(j)}$ with the sequence $w^{(j)}$ and repeat the same process for the 1st and i -th coordinate, for $i = 3, \dots, d$. Since each step only zeroes the 1st coordinate of some vectors and decreases the other coordinates, the final resulting sequence $w^{(j)}$, when compared to the initial sequence $v^{(j)}$ satisfies:

$$\sum_{j \in J} |v_1^{(j)} - w_1^{(j)}| \leq (d-1) \cdot \frac{\varepsilon}{d} \sum_{j \in J} v_1^{(j)}$$

$$\sum_{j \in J} |v_i^{(j)} - w_i^{(j)}| \leq (\frac{\varepsilon}{2} + \frac{\varepsilon}{2}) \sum_{j \in J} v_i^{(j)} \quad \text{for } i = 2, \dots, d.$$

Each vector $w^{(j)}$ either has its 1st coordinate zeroed, or all its other coordinates are determined as $w_1^{(j)}$ times one of $2 + K$ possible ratios. Among vectors with $w_1^{(j)} \neq 0$ there are thus at most $(2 + K)^{d-1}$ different vectors, up to rescaling. The vectors with $w_1^{(j)} = 0$ can be inductively reduced as $(d-1)$ -dimensional vectors to $w^{(j)'} containing $k(d-1, \frac{\varepsilon}{3})$ distinct vectors up to rescaling (where $k(d-1, \frac{\varepsilon}{3})$ is the constant k obtained by inductive assumption for $d-1$ and $\frac{\varepsilon}{3}$) and satisfying$

$$\sum_{j \in J} |w_i^{(j)} - w_i^{(j)'}| \leq \frac{\varepsilon}{3} \sum_{j \in J} w_i^{(j)} \leq \frac{\varepsilon}{3} \sum_{j \in J} v_i^{(j)}.$$

Altogether, the difference is bounded by $\frac{2\varepsilon}{3} + \frac{\varepsilon}{3} = \varepsilon$ and the number of distinct vectors up to rescaling is bounded by $(2 + K)^{d-1} + k(d-1, \frac{\varepsilon}{3})$. \square

Lemma 5.5. *For a fixed signature σ and $d \in \mathbb{N}$, the class of σ -structures with maximum connected component size at most d is size-pliable.*

Proof. We simply present each component \mathbb{A}_i of \mathbb{A} as a vector encoding the value of all tuples $(f, \mathbf{x}) \in \text{tup}(\mathbb{A}_i)$. The dimension of such a vector, for a component of size d , is $d' := \sum_{f \in \sigma} d^{\text{ar}(f)}$. Smaller components can be treated as components of size d by adding dummy vertices and tuples.

For any $\varepsilon > 0$, let $\varepsilon' := \frac{\varepsilon/C_\sigma}{1+\varepsilon/C_\sigma}$, where $C_\sigma = \max_{f \in \sigma} \text{ar}(f)^{\text{ar}(f)}$. The previous lemma guarantees the existence of a number $k = k(\varepsilon', d')$ such that for every σ -structure \mathbb{A} with n components of size at most d , the corresponding vectors $v^{(1)}, \dots, v^{(n)}$ are approximated by vectors $w^{(1)}, \dots, w^{(n)}$ such that there are at most k distinct vectors up to rescaling and such that, for $i = 1 \dots n$,

$$\sum_{j=1 \dots n} |v_i^{(j)} - w_i^{(j)}| \leq \varepsilon' \sum_{j=1 \dots n} v_i^{(j)}.$$

These vectors encode a σ -structure \mathbb{B} with only at most k distinct components up to rescaling, all of size at most d , which is hence (by Observation 5.3) equivalent to a σ -structure \mathbb{B}' bounded in size by $k \cdot d$. Moreover, the guarantee on ε' allows us to bound edit distance as follows:

$$\sum_{j=1 \dots n} |v_i^{(j)} - w_i^{(j)}| \leq \varepsilon' \sum_{j=1 \dots n} v_i^{(j)} \leq \varepsilon' \sum_{j=1 \dots n} (\min(v_i^{(j)}, w_i^{(j)}) + |v_i^{(j)} - w_i^{(j)}|)$$

(for $i = 1 \dots d'$), hence the edit distance (as defined in Section 2.4) is

$$d_1(\mathbb{A}, \mathbb{B}) \leq \max_{i=1 \dots d'} \frac{\sum_{j=1 \dots n} |v_i^{(j)} - w_i^{(j)}|}{\min(\sum_{j=1 \dots n} v_i^{(j)}, \sum_{j=1 \dots n} w_i^{(j)})} \leq \frac{\varepsilon'}{1 - \varepsilon'} = \frac{\varepsilon}{C_\sigma}.$$

By Lemma 2.11, $d_{\text{opt}}(\mathbb{A}, \mathbb{B}') \leq d_{\text{opt}}(\mathbb{A}, \mathbb{B}) + d_{\text{opt}}(\mathbb{B}, \mathbb{B}') = d_{\text{opt}}(\mathbb{A}, \mathbb{B}) \leq C_\sigma \cdot d_1(\mathbb{A}, \mathbb{B}) \leq \varepsilon$. (While Lemma 2.11 assumes the structures to be loopless, this can be ensured by replacing tuples with repeated elements like $(f, (x_1, x_1, x_2))$, say, with $(f', (x_1, x_2))$ for a new symbol f'). \square

By Observation 5.1 (transitivity of pliability), we conclude that for a fixed signature σ , if a class of σ -structures \mathcal{A} is cc-pliable then it is also size-pliable. Thus, we have shown equivalence of size-pliability and cc-pliability (for structures of bounded signatures).

5.3 From treedepth-pliability to size-pliability

In order to finish the proof of Theorem 1.5, it remains to show the equivalence of td- and size-pliability. We do this by extending the above proof for cc- and size-pliability.

One of the main reasons for which treedepth is useful (and easier to work with than, say, treewidth) is that the only way for a graph of small treedepth to be large is to have many repeating parts, like in a large star graph (see e.g. [74, Theorem 3.1]). This implies that in a class of graphs of bounded treedepth, homomorphic cores have bounded size. This does not extend to weighted graphs or structures in general, but we can approximate the weights or values as before.

Lemma 5.6. *For a fixed signature σ and $d \in \mathbb{N}$, the class of σ -structures $\{\mathbb{A} : \text{td}(\mathbb{A}) \leq d\}$ is size-pliable.*

Proof. We prove by induction on d that the statement holds for each signature σ . It suffices to prove the statement for connected σ -structures of treedepth at most d . Indeed, this implies that disconnected σ -structures of treedepth at most d are cc-pliable, which we already know implies size-pliability.

For $d = 1$, each component of the Gaifman graph is a single vertex and we are done. So let $d > 1$ and assume that for each signature σ and each $\varepsilon > 0$, there is a $k = k(d - 1, \sigma, \varepsilon)$ such that every σ -structure with treedepth $\leq d - 1$ has a σ -structure of size $\leq k$ at opt-distance at most ε . Let σ be a signature and \mathbb{A} a σ -structure of treedepth d . Let G be the Gaifman graph of \mathbb{A} . Since it is connected, we can find a vertex $v \in V(G) = A$ such that $\text{td}(G - v) = d - 1$.

We now define a new signature σ' and a σ' -structure $\text{pack}(\mathbb{A})$ whose Gaifman graph will be $G - v$, but will contain all the information about \mathbb{A} . Let $\sigma' = \{(f, I) : f \in \sigma, I \subseteq \{1, \dots, \text{ar}(f)\}\}$ and $\text{ar}((f, I)) := \text{ar}(f) - |I|$, for $(f, I) \in \sigma'$. For $\mathbf{x} \in (A - v)^{\text{ar}((f, I))}$, let $(f, I)^{\text{pack}(\mathbb{A})}(\mathbf{x}) := f^{\mathbb{A}}(\mathbf{x}')$ where $\mathbf{x}' \in A^{\text{ar}(f)}$ is the tuple obtained from \mathbf{x} by introducing v at positions I . Note that σ' is bounded: $|\sigma'| = \sum_{f \in \sigma} 2^{\text{ar}(f)}$.

The σ' -structure $\text{pack}(\mathbb{A})$ has treedepth $d - 1$, so by inductive assumption there is a σ' -structure \mathbb{B} at opt-distance at most ε with size at most $k = k(d - 1, \sigma', \varepsilon)$. We define $\text{unpack}(\mathbb{B})$ to be the σ -structure with domain $B \cup \{v\}$ and $f^{\text{unpack}(\mathbb{B})}(\mathbf{x}) := (f, I)^{\mathbb{B}}(\mathbf{x}')$, where I is the set of positions in \mathbf{x} containing v and \mathbf{x}' is the tuple obtained by removing these positions. It is straightforward to check that $\text{unpack}(\text{pack}(\mathbb{A}))$ is equal to \mathbb{A} and that for any σ' -structures \mathbb{A}', \mathbb{B}' we have $d_{\text{opt}}(\text{unpack}(\mathbb{A}'), \text{unpack}(\mathbb{B}')) \leq d_{\text{opt}}(\mathbb{A}', \mathbb{B}')$, hence

$$d_{\text{opt}}(\mathbb{A}, \text{unpack}(\mathbb{B})) = d_{\text{opt}}(\text{unpack}(\text{pack}(\mathbb{A})), \text{unpack}(\mathbb{B})) \leq d_{\text{opt}}(\text{pack}(\mathbb{A}), \mathbb{B}) \leq \varepsilon.$$

Hence $\text{unpack}(\mathbb{B})$ is a σ -structure at opt-distance $\leq \varepsilon$ from \mathbb{A} of size $\leq k + 1$. \square

By Observation 5.1 (transitivity of pliability), this shows that, for a class of σ -structures, td-pliability implies size-pliability. Since $\text{td} \leq \text{size}$, Observation 2.9 shows that size-pliability implies td-pliability, thus concluding the proof of Theorem 1.5:

Corollary 5.7. *A class of σ -structures is td-pliable if and only if it is size-pliable.*

6 Hyperfinite classes are fractionally fragile: proof of Theorem 1.6

Recall that class of graphs is *hyperfinite* if for every $\varepsilon > 0$ there is a $k \in \mathbb{N}$ such that every graph in the class can be turned into a graph with connected components of size at most k by removing an at-most- ε fraction of all edges.¹⁰ A class of graphs is *monotone* if it is closed under taking subgraphs. In this section, we prove the following result.

Theorem (Theorem 1.6 restated). *Let \mathcal{G} be a monotone class of graphs. The following are equivalent:*

- \mathcal{G} is hyperfinite;
- \mathcal{G} is fractionally-tw-fragile and has bounded degree;
- \mathcal{G} is fractionally-cc-fragile;
- $\mathcal{A}_{\mathcal{G}}^{(r)}$ is cc-pliable for any $r \geq 2$.

The last two bullets are shown equivalent by Lemma 4.7; the middle two bullets were shown equivalent by Dvořák [29, Observation 15, Corollary 20]. It remains to prove their equivalence with the first bullet point.

Lemma 6.1. *Let \mathcal{G} be a monotone class of graphs. \mathcal{G} is hyperfinite if and only if it is fractionally-cc-fragile.*

Proof. Hyperfiniteness of a monotone class \mathcal{G} is equivalent to hyperfiniteness of 0-1-edge weighted graphs in \mathcal{G} :

$$\forall \varepsilon > 0 \exists k \forall G \in \mathcal{G} \forall w: E(G) \rightarrow \{0,1\} \exists F \subseteq E(G) \ w(F) \leq \varepsilon w(E(G)) \text{ and } \text{cc}(G - F) \leq k.$$

¹⁰In other work, the definition of *hyperfinite* often considers the number of removed edges divided by the total number of vertices. However, they only deal with bounded degree graphs (in which the number of edges is linear in the number of vertices), which makes the two definitions equivalent.

Hence it is trivially implied by the edge version of fractional-cc-fragility (which allows arbitrary nonnegative weights) in Lemma 4.4. It remains to show the other direction.

By definition of hyperfiniteness, for every $\varepsilon > 0$ there is a $k = k(\varepsilon)$ such that for all graphs $G \in \mathcal{G}$, one can remove a set of edges F with $|F| \leq \varepsilon|E(G)|$ so that $\text{cc}(G - F) \leq k$. Observe that graphs in \mathcal{G} have degree bounded by $\Delta := 2k(\frac{1}{2})$; otherwise, a graph with degree $\geq 2k(\frac{1}{2}) + 1$ would contain a star with that many edges as a subgraph and removing half of these edges always leaves a component with at least $k(\frac{1}{2}) + 1$ edges and vertices.

We aim to show that

$$\forall \varepsilon > 0 \exists k \forall G \in \mathcal{G} \forall w: E(G) \rightarrow \mathbb{Q}_{\geq 0} \exists F \subseteq E(G) \quad w(F) \leq \varepsilon w(E(G)) \text{ and } \text{cc}(G - F) \leq k.$$

For $\varepsilon > 0$, let ε' be chosen later and let $k' = k(\varepsilon')$. Let $G \in \mathcal{G}$ and $w: E(G) \rightarrow \mathbb{Q}_{\geq 0}$. We want to find a set $F \subseteq E(G)$ such that $w(F) \leq \varepsilon w(E(G))$ and $\text{cc}(G - F) \leq k'$. Note that our task would be trivial if the weights of all edges were within a constant factor α of each other: just set $\varepsilon' = \frac{\varepsilon}{\alpha}$, find $F \subseteq E(G)$ such that $|F| \leq \varepsilon'|E(G)|$ and $\text{cc}(G - F) \leq k'$ and conclude that $w(F) \leq \alpha \varepsilon' w(E(G)) = \varepsilon w(E(G))$.

In general, let us partition the edges of G into buckets depending on their weight: for $i \in \mathbb{Z}$, let $B_i := \{e \in E(G) \mid (\frac{\varepsilon}{6\Delta})^i \geq w(e) > (\frac{\varepsilon}{6\Delta})^{i+1}\}$ (edges with weight zero can be removed without loss of generality).¹¹ For $L := \lceil \frac{3}{\varepsilon} \rceil$, we will remove every L -th bucket from G . That is, for $j \in \{0, \dots, L-1\}$, let $B'_j := \bigcup_{i \in \mathbb{Z}} B_{iL+j}$. Let $j^* \in \{0, \dots, L-1\}$ be such that $w(B'_{j^*})$ is minimum; since $B'_0 \cup \dots \cup B'_{L-1}$ is a partition of $E(G)$, $w(B'_{j^*}) \leq \frac{1}{L}w(E(G)) \leq \frac{\varepsilon}{3}w(E(G))$. We can thus remove the edges B'_{j^*} from G . Since this removes every L -th bucket, the remaining edges are partitioned into blocks $C_i := B_{iL+j^*+1} \cup \dots \cup B_{iL+j^*+L-1}$ of $L-1$ buckets for $i \in \mathbb{Z}$. Each block contains weights within a constant factor of each other: $\min\{w(e) : e \in C_i\} \geq (\frac{\varepsilon}{6\Delta})^{L-1} \cdot \max\{w(e) : e \in C_i\}$. Moreover, since there is a gap of one bucket in between one block and the next, $\max\{w(e) : e \in C_{i+1}\} < \frac{\varepsilon}{6\Delta} \cdot \min\{w(e) : e \in C_i\}$.

The latter property allows us to disconnect the blocks from each other. Indeed, for each C_i with increasing i (starting from the smallest i such that C_i is non-empty), we shall remove all remaining edges on the boundary of C_i :

$$F_i := \{e : e \in C_j \text{ for some } j > i, e \text{ shares a vertex with some } e' \in C_i\}.$$

Since $|F_i| \leq 2\Delta|C_i|$ and $\max\{w(e) : e \in F_i\} < \frac{\varepsilon}{6\Delta} \cdot \min\{w(e) : e \in C_i\}$, we have $w(F_i) < \frac{\varepsilon}{3}w(C_i)$. In total, for $F := \bigcup_{i \in \mathbb{Z}} F_i$, we have $w(F) < \frac{\varepsilon}{3}w(E(G))$. For all $i \in \mathbb{Z}$, the edges of $C_i - F$ are disjoint from edges of $C_j - F$ for all $j > i$. Since the sets $C_i - F$ partition edges of $G - B'_{j^*} - F$, this means that every connected component of $G - B'_{j^*} - F$ is contained in one of the edge sets $C_i - F$.

Finally, since $\min\{w(e) : e \in C_i - F\} \geq \alpha \cdot \max\{w(e) : e \in C_i - F\}$ for $\alpha := (\frac{\varepsilon}{3\Delta})^{L-1}$, we have reduced our problem to the trivial case when weights are all within a constant factor of each other. That is, let $\varepsilon' := \frac{\alpha\varepsilon}{3}$. For each $i \in \mathbb{Z}$, let G_i be the subgraph of $G - B'_{j^*} - F$ formed from connected components contained in $C_i - F$. By assumption, there is a set $F'_i \subseteq E(G_i)$ such that $|F'_i| \leq \varepsilon'|E(G_i)|$ and $\text{cc}(G_i - F'_i) \leq k(\varepsilon') = k'$. Then $w(F'_i) \leq \max\{w(e) : e \in E(G_i)\} \cdot |F'_i| \leq \frac{1}{\alpha} \cdot \min\{w(e) : e \in E(G_i)\} \cdot \varepsilon' \cdot |E(G_i)| \leq \frac{\varepsilon}{\alpha}w(E(G_i)) = \frac{\varepsilon}{3}w(E(G_i))$. In total, for $F' := \bigcup_{i \in \mathbb{Z}} F'_i$ we have $w(B'_{j^*} \cup F \cup F') \leq (\frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3})w(E(G))$ and $\text{cc}(G - B'_{j^*} - F - F') \leq k'$. \square

¹¹Note that all but a finite number of B_i 's will be empty.

7 Dense graphs are pliable: proof of Theorem 1.3

Our goal is to prove Theorem 1.3.

Theorem (Theorem 1.3 restated). *Let $c > 0$ and let \mathcal{A} be a class of graphs with at least cn^2 edges. Then \mathcal{A} is tw-pliable. Consequently, $\text{Max-Hom}(\mathcal{A}, -)$ admits a PTAS.*

In order to do so, we prove the following result.

Theorem 7.1. *Let $c > 0$. The class of (unweighted, undirected) graphs with at least cn^2 edges is size-pliable.*

Theorem 7.1 implies Theorem 1.3. Indeed, since $\text{tw} \leq \text{size}$, by Observation 2.9 we have size-pliability implies tw-pliability. By Theorem 1.1, tw-pliability implies a PTAS.

We start with simple examples of dense graphs. Observe that large cliques can be arbitrarily well approximated by cliques of constant size $\lceil \frac{2}{\varepsilon} \rceil$ (up to normalising total edge weights).

Example 7.2. Let $0 < \varepsilon < 1$ and let $n, k \geq \frac{2}{\varepsilon}$. Then $d_{\text{opt}}(K_n, \lambda K_k) \leq \varepsilon$, for $\lambda = \binom{n}{2} / \binom{k}{2}$.

Proof. For $n, k \geq 2$, define an overcast ω by taking a random function $V(K_n) \rightarrow V(K_k)$ (each vertex is placed independently uniformly at random). Then for each $e \in E(K_k)$,

$$\mathbb{E}_{g \sim \omega} |g^{-1}(e)| = \sum_{e' \in E(K_n)} \mathbb{E}_{g \sim \omega} [g(e') = e] = \binom{n}{2} \frac{2}{k^2} = \lambda \cdot \binom{k}{2} \cdot \frac{2}{k^2} = (1 - \frac{1}{k})\lambda.$$

Therefore $K_n \succeq (1 - \frac{1}{k})\lambda K_k$. Symmetrically $\lambda K_k \succeq (1 - \frac{1}{n})K_n$. Since $1 - x \geq e^{-2x}$ for $0 \leq x \leq \frac{1}{2}$, this means $d_{\text{opt}}(K_n, \lambda K_k) \leq \frac{2}{\min(n, k)}$. Consequently if $n, k \geq \frac{2}{\varepsilon}$, then $d_{\text{opt}}(K_n, \lambda K_k) \leq \varepsilon$. \square

In particular, this means the class \mathcal{A} consisting of all clique graphs is size-pliable. This corresponds to an easy PTAS for graph $\text{Max-Hom}(\mathcal{A}, -)$: the maximum graph homomorphism from K_n to G is well approximated by finding the maximum graph homomorphism from a constant size K_k to G and mapping K_n randomly to the resulting $\leq k$ vertices in G . The situation is very different for Densest Subgraph problems, because they disallow choosing two equal vertices in G (see Observation 8.2).

As another important example, consider Erdős-Rényi random graphs $G(n, p)$ (for constant $p \in (0, 1)$; each pair in $\binom{n}{2}$ becomes an edge independently with probability p). Any two such graphs are similar to each other (and in fact to pK_n , as well as to λK_k for constant k and suitable λ); more precisely, we have:

Example 7.3. Let $p, \varepsilon > 0$ be constants. Let G_1, G_2 be independent Erdős-Rényi random graphs $G(n, p)$. Then $\Pr[d_{\text{opt}}(G_1, G_2) < \varepsilon] \rightarrow 1$ as $n \rightarrow \infty$.

Proof sketch. Let k be a sufficiently large constant depending on ε only. It is sufficient to prove that $\Pr[d_{\text{opt}}(G_1, \lambda K_k) < \frac{\varepsilon}{2}] \rightarrow 1$ as $n \rightarrow \infty$. The rescaling factor here is $\lambda := p \binom{n}{2} / \binom{k}{2}$. The number of edges of $G(n, p)$ is concentrated around $p \binom{n}{2}$, so just as before a random function gives $G(n, p) \succeq (1 - \frac{1}{k})\lambda K_k \geq e^{-\varepsilon/2} \lambda K_k$ with high probability (tending to 1 as $n \rightarrow \infty$).

For the other direction, we use the fact that the number of k -cliques in $G(n, p)$ is concentrated around the mean $\binom{n}{k} p^{\binom{k}{2}}$ and, more strongly, the number of k -cliques containing any given edge of $G(n, p)$ (conditioned on it being an edge) is concentrated around the mean

$\binom{n-2}{k-2}p^{\binom{k}{2}-1}$. The concentration is good enough that with high probability, *every* edge of $G(n, p)$ is contained in $(1 \pm \frac{\varepsilon}{4})\binom{n-2}{k-2}p^{\binom{k}{2}-1}$ k -cliques (see e.g. [84]). Thus if we take ω by mapping λK_k injectively to a random k -clique in $G(n, p)$, then w.h.p. for each edge e of $G(n, p)$ we have

$$\mathbb{E}_{g \sim \omega} |g^{-1}(e)| \geq (1 - \frac{\varepsilon}{4}) \binom{n-2}{k-2} p^{\binom{k}{2}-1} / \binom{n}{k} p^{\binom{k}{2}} = (1 - \frac{\varepsilon}{4}) \frac{k(k-1)}{n(n-1)} p^{-1} = (1 - \frac{\varepsilon}{4}) \lambda^{-1}.$$

Thus $\lambda K_k \succeq e^{-\varepsilon/2} G_1$ and consequently $d_{\text{opt}}(G_1, \lambda K_k) \leq \frac{\varepsilon}{2}$ w.h.p. \square

To show Theorem 1.3, we extend the above informal proof to any class of dense graphs. This is possible because of *Szemerédi's regularity lemma* [85], which, very roughly speaking, guarantees that all such graphs are random-like. This allows to provide similar bounds on the number of k -cliques containing any given edge, a fact known as the *extension lemma*, though we prove a variant that is somewhat tighter than usual.

Remark 7.4. Note that the above proof sketch does not work for random tournaments (orientations of cliques): if we try to approximate them by the small graph $\frac{1}{2} \overleftrightarrow{K}_k$ (each arc taken with weight $\frac{1}{2}$), then every overcast from it to a tournament will always lose at least half of the total weight. If instead we tried to take a small random tournament, no overcast to it from the big random tournament would work. Indeed, Lemma 2.19 in Section 2.6 shows the class of tournaments is not pliable (neither are “random tournaments”, i.e., the proof can be adapted to show that any class which contains a random tournament with constant probability cannot be pliable) and in fact the problem $\text{Max-Hom}(\mathcal{A}, -)$ for the class of tournaments \mathcal{A} is hard to approximate, as we show in Lemma 8.4 in Section 8. This is why, even though variants of the regularity lemma exist for directed graphs and even more general structures, we limit our discussion to undirected graphs (the proofs do extend to $[0, 1]$ -weighted undirected graphs, however).

In the rest of this section, we will prove Theorem 1.3. While we only prove this for unweighted graphs, it will be notationally convenient to treat them as $\{0, 1\}$ -weighted graphs, with $w_G(u, v) := [uv \in E(G)]$. For sets $U, V \subseteq V(G)$, we denote by $w_G(U, V) := \sum_{u \in U} \sum_{v \in V} w_G(u, v)$ the number of edges between U and V (or their total weight). The regularity lemma states that every graph can be partitioned into a bounded number of parts so that the bipartite graph between every two parts is random-like in the following strong sense:

Definition 7.5. A bipartite graph $G = (V_1, V_2, E)$ of density $d := \frac{w_G(V_1, V_2)}{|V_1||V_2|}$ is ε -homogeneous if for all $W_1 \subseteq V_1, W_2 \subseteq V_2$,

$$w_G(W_1, W_2) = d|W_1||W_2| \pm \varepsilon|V_1||V_2|.$$

For an n -vertex graph G and an integer k , an ε -regular k -partition of G is a partition V_1, \dots, V_k of $V(G)$ such that $|V_i| \in \{\lfloor \frac{n}{k} \rfloor, \lceil \frac{n}{k} \rceil\}$ for $i \in [k]$ and the bipartite graph $(V_i, V_j, E(G) \cap V_i \times V_j)$ is ε -homogeneous¹², for all $ij \in \binom{[k]}{2}$.

We use the following strong version of Szemerédi's regularity lemma (see Theorem 2.2 in [79], Lemma 5.2. in [67], or Chapter 9 in [66] for a detailed discussion).

¹²The usual statement of the regularity lemma replaces ε -homogeneity (with additive error) by a notion called ε -regularity (with relative error, but holding only for $|W_i| \geq \varepsilon|V_i|$). The two are however easily shown to be equivalent, up to a change from ε to $\varepsilon^{1/3}$.

Theorem 7.6 (Regularity Lemma). *For every $\varepsilon_1 > 0$ and every non-decreasing $f: \mathbb{N} \rightarrow \mathbb{N}$, there is an integer k such that for every sufficiently large graph G , one can add/remove $\varepsilon_1|V(G)|^2$ edges to obtain a graph which admits a $\frac{1}{f(k')}$ -regular k' -partition for some $\frac{1}{\varepsilon_1} \leq k' \leq k$.*

Another way to view this is to define, for a partition $\mathcal{P} = (V_1, \dots, V_k)$ of a graph G , the *quotient graph* $G_{/\mathcal{P}}$ as the weighted graph with vertex set $[k]$ and weights $w_{G_{/\mathcal{P}}}(i, j) := w_G(V_i, V_j)$ for $(i, j) \in [k]^2$. The quotient graph for an ε -regular partition is then a graph of bounded size that is *close* to the original graph: the notion of closeness arising from the definition of ε -homogeneity is known as *cut distance* (see Chapter 8 in [66]), but later we show the same holds for opt-distance:

Theorem 7.7. *Let G be a graph with density $c := \frac{|E(G)|}{n^2}$. For $0 < \varepsilon_0 < 1$, suppose G has an ε -regular k -partition $\mathcal{P} = (V_1, \dots, V_k)$ with $\frac{1}{k} \leq \frac{c}{10} \frac{\varepsilon_0}{1+\varepsilon_0}$ and $\varepsilon \leq \left(\frac{1}{k}\right)^{8k^2}$. Then $d_{\text{opt}}(G, G_{/\mathcal{P}}) \leq \varepsilon_0$.*

With this view it is easy to see that classes of dense graphs are pliable. Formally:

Proof of Theorem 1.3 assuming Theorem 7.7. Let \mathcal{A} be a class of graphs with $\geq cn^2$ edges. We want to show that for every $\varepsilon_0 > 0$ there is a k such that every $G \in \mathcal{A}$ has a weighted graph H of size at most k with $d_{\text{opt}}(G, H) \leq \varepsilon_0$.

For $\varepsilon_0 > 0$, let $\varepsilon_1 := \frac{c}{10} \cdot \frac{\varepsilon_0/2}{1+\varepsilon_0/2}$. Note that we can assume that all sufficiently large graphs $G \in \mathcal{A}$ have no loops: if $|V(G)| \geq \frac{1}{\varepsilon_1}$, then the number of loops is at most $|V(G)| \leq \frac{1}{c|V(G)|}|E(G)| \leq \frac{\varepsilon_1}{c}|E(G)|$. Hence by removing them we obtain a graph G' such that $d_1(G, G') \leq \frac{|E(G)| - |E(G')|}{\min(|E(G)|, |E(G')|)} \leq \frac{\varepsilon_1/c}{1-\varepsilon_1/c} \leq \varepsilon_0/20$. By Lemma 2.11, $d_{\text{opt}}(G, G') \leq \varepsilon_0/5$ (the direction $G \succeq G'$ is trivial, while the other direction only requires G' to have no loops).

Let $f(k) := k^{8k^2}$. By the Regularity Lemma (Theorem 7.6), there is an integer $k \geq \frac{1}{\varepsilon_1}$ such that for every graph G of size $> k$, one can add/remove $\varepsilon_1|V(G)|^2$ edges to obtain a graph H which admits a $\frac{1}{f(k')}$ -regular k' -partition \mathcal{P} for some $\frac{1}{\varepsilon_1} \leq k' \leq k$. If $G \in \mathcal{A}$, then G has at least cn^2 edges, so $d_1(G, H) \leq \frac{\varepsilon_1}{c-\varepsilon_1} \leq \varepsilon_0/20$. Since we can assume that G is loopless, by Lemma 2.11, $d_{\text{opt}}(G, H) \leq \varepsilon_0/5$. By Theorem 7.7, $d_{\text{opt}}(H, H_{/\mathcal{P}}) \leq \varepsilon_0/2$. Hence $H_{/\mathcal{P}}$ is the graph of size at most k we seek, at opt-distance at most ε_0 from G . \square

The strategy of the proof of Theorem 7.7 is very similar to Example 7.3. One direction is trivial: an overcast from G to $G_{/\mathcal{P}}$ is given simply by deterministically mapping all of V_i to i , for $i \in [k]$. For the other direction, we will take a subgraph F of $G_{/\mathcal{P}}$ obtained by removing edges of small weight (keeping F close to $G_{/\mathcal{P}}$) and removing weights, and then map $G_{/\mathcal{P}}$ to a random copy of F in G . We need to estimate the number of such copies (this is known as the *counting lemma*) and, more generally, the number of such copies containing any given edge of G (the *extension lemma*). Both are standard lemmas in the theory of dense graph limits, in particular our proof of the counting lemma mimics Lemma 10.22 in [66]. However, we will prove a version of the extension lemma with somewhat tighter bounds than usual (depending on all $\binom{k}{2}$ edge densities between parts of the regularity partition).

For a graph F on vertex set $[k] := \{1, \dots, k\}$, we will treat F as a subset of $\binom{[k]}{2}$. For a partition $\mathcal{P} = (V_1, \dots, V_k)$ of a graph G , a \mathcal{P} -map is a function $g: [k] \rightarrow V(G)$ such that $g(i) \in V_i$ for all $i \in [k]$. We denote $\text{hom}_g(F, G) := \prod_{ij \in F} w_G(g(i), g(j))$; for $\{0, 1\}$ -weighted graphs, this is equal to 1 if g is a homomorphism from F to G and 0 otherwise.

Let us first observe two consequences of ε -homogeneity. First, the notion can be extended from subsets $W_1 \subseteq V_1$ to any function $f: V_1 \rightarrow [0, 1]$ simply by considering subsets where the function takes at least a given value (here $\|f\|_1 := \sum_x f(x)$):

Observation 7.8. *Let $G = (V_1, V_2, E)$ be ε -homogeneous of density d . Then for every $f: V_1 \rightarrow [0, 1]$ and $g: V_2 \rightarrow [0, 1]$,*

$$\sum_{x_1 \in V_1} \sum_{x_2 \in V_2} f(x_1)g(x_2)w_G(x_1, x_2) = d\|f\|_1\|g\|_1 \pm \varepsilon|V_1||V_2|.$$

Proof. For $y \in [0, 1]$, let $V_1^y := \{x \in V_1: f(x) \geq y\}$ and define V_2^y analogously for g . Notice that $f(x) = \int_0^1 [y \leq f(x)] dy = \int_0^1 [x \in V_1^y] dy$. Hence

$$\begin{aligned} \sum_{x_1 \in V_1} \sum_{x_2 \in V_2} f(x_1)g(x_2)w_G(x_1, x_2) &= \\ \int_0^1 \int_0^1 \sum_{x_1 \in V_1} \sum_{x_2 \in V_2} [x_1 \in V_1^{y_1}][x_2 \in V_2^{y_2}]w_G(x_1, x_2) dy_1 dy_2 &= \\ \int_0^1 \int_0^1 w_G(V_1^{y_1}, V_2^{y_2}) dy_1 dy_2 &= \\ \int_0^1 \int_0^1 (d|V_1^{y_1}||V_2^{y_2}| \pm \varepsilon|V_1||V_2|) dy_1 dy_2 &= \\ d \cdot \left(\int_0^1 |V_1^{y_1}| dy_1 \right) \left(\int_0^1 |V_2^{y_2}| dy_2 \right) \pm \varepsilon|V_1||V_2|. \end{aligned}$$

Since $\int_0^1 |V_1^y| dy = \int_0^1 \sum_{x \in V_1} [y \leq f(x)] dy = \sum_{x \in V_1} f(x) = \|f\|_1$ and analogously for g , the claim follows. \square

Second, while we cannot say much about any one fixed vertex, we can make similar approximations for most vertices:

Observation 7.9. *Let $G = (V_1, V_2, E)$ be ε -homogeneous with density d . For every $g: V_2 \rightarrow [0, 1]$, there are at least $(1 - 2\sqrt{\varepsilon})|V_1|$ vertices x_1 in V_1 such that $\sum_{x_2} g(x_2)w_G(x_1, x_2) = d\|g\|_1 \pm \sqrt{\varepsilon}|V_2|$.*

Proof. Let W_1^- be the set of those x_1 in V_1 for which the sum is too small:

$$\sum_{x_2} g(x_2)w_G(x_1, x_2) < d\|g\|_1 - \sqrt{\varepsilon}|V_2|.$$

Let $f: V_1 \rightarrow [0, 1]$ be the characteristic function of W_1^- . Then

$$\sum_{x_1} \sum_{x_2} f(x_1)g(x_2)w_G(x_1, x_2) < \|f\|_1 \cdot (d\|g\|_1 - \sqrt{\varepsilon}|V_2|).$$

By Observation 7.8, this implies $\|f\|_1 \cdot \sqrt{\varepsilon}|V_2| < \varepsilon|V_1||V_2|$. Hence $|W_1^-| = \|f\|_1 < \sqrt{\varepsilon}|V_1|$. We can define and bound W_1^+ analogously. Then $V_1 \setminus (W_1^- \cup W_1^+)$ is a set of size at least $(1 - 2\sqrt{\varepsilon})|V_1|$ as claimed. \square

The counting lemma says that the number of \mathcal{P} -maps that are homomorphisms from F to G is close to what one would expect in a purely random graph with the same densities. Note that the number of all \mathcal{P} -maps $g: [k] \rightarrow V(G)$ is exactly $\prod_{i \in [k]} |V_i|$.

Lemma 7.10 (Counting Lemma). *Let $\mathcal{P} = (V_1, \dots, V_k)$ be an ε -regular k -partition of an n -vertex graph G . Let $d_{ij} := \frac{w_G(V_i, V_j)}{|V_i||V_j|}$. For each $F \subseteq \binom{[k]}{2}$,*

$$\sum_g \text{hom}_g(F, G) = \left(\prod_{i \in [k]} |V_i| \right) \left(\prod_{ij \in F} d_{ij} \pm \varepsilon |F| \right),$$

where the sum is over all \mathcal{P} -maps $g: [k] \rightarrow V(G)$.

Proof. Let us write $\sum_{(x_i)_{i \in [k]}}$ as a shorthand for $\sum_{x_1 \in V_1} \cdots \sum_{x_k \in V_k}$. We wish to approximate

$$\sum_{(x_i)_{i \in [k]}} \prod_{ij \in F} w_G(x_i, x_j).$$

We do so by replacing each factor $w_G(x_i, x_j)$ by its average value d_{ij} , one by one. That is, we prove for all subsets $F' \subseteq F$ by induction that

$$\sum_{(x_i)_{i \in [k]}} \prod_{ij \in F} w_G(x_i, x_j) = \sum_{(x_i)_{i \in [k]}} \prod_{ij \in F - F'} w_G(x_i, x_j) \prod_{ij \in F'} d_{ij} \pm |F'| \cdot \varepsilon \prod_{i \in [k]} |V_i|. \quad (*)$$

Clearly this is true initially for $F' = \emptyset$ and eventually by reaching $F' = F$ we will have proved that

$$\sum_{(x_i)_{i \in [k]}} \prod_{ij \in F} w_G(x_i, x_j) = \sum_{(x_i)_{i \in [k]}} \prod_{ij \in F} d_{ij} \pm |F| \cdot \varepsilon \prod_{i \in [k]} |V_i|;$$

which proves the claim, as $\sum_{(x_i)_{i \in [k]}} 1 = \prod_{i \in [k]} |V_i|$.

To prove the induction step, suppose $(*)$ holds for some $F' \subset F$ and let $ab \in F - F'$. Let $w'_G(x_i, x_j)$ denote $w_G(x_i, x_j)$ if $ij \notin F'$ and d_{ij} otherwise. Then the left-hand-side in $(*)$ is

$$\begin{aligned} & \sum_{(x_i)_{i \in [k]}} \prod_{ij \in F} w'_G(x_i, x_j) = \\ & \sum_{(x_i)_{i \in [k] - a - b}} h \sum_{x_a} \sum_{x_b} f(x_a) g(x_b) w'_G(x_a, x_b), \end{aligned}$$

where for any fixed choice of $(x_i)_{i \in [k] - a - b}$, we let

$$h := \prod_{\substack{ij \in F - F' \\ i \neq a, j \neq b}} w'_G(x_i, x_j), \quad f(x_a) := \prod_{\substack{ij \in F - F' \\ i = a, j \neq b}} w'_G(x_a, x_j), \quad g(x_b) := \prod_{\substack{ij \in F - F' \\ i \neq a, j = b}} w'_G(x_i, x_b).$$

Since $h, f, g \leq 1$, the claim then follows from Observation 7.8: replacing $w_G(x_a, x_b)$ with d_{ab} adds an error of at most $\sum_{(x_i)_{i \in [k] - a - b}} h \cdot \varepsilon |V_a| |V_b| \leq \varepsilon |V_a| |V_b| \cdot \sum_{(x_i)_{i \in [k] - a - b}} 1 = \varepsilon \prod_{i \in [k]} |V_i|$. \square

Lemma 7.11 (Extension Lemma). *Let $\mathcal{P} = (V_1, \dots, V_k)$ be an ε -regular k -partition of an n -vertex graph G . Let $d_{ij} := \frac{w_G(V_i, V_j)}{|V_i||V_j|}$. For each $F \subseteq \binom{[k]}{2}$ and each $ab \in F$, all but $2k\sqrt{\varepsilon}|V_a||V_b|$ edges $x_a x_b \in V_a \times V_b$ satisfy*

$$\sum_g \text{hom}_g(F, G) = \left(\prod_{i \in [k] - a - b} |V_i| \right) \cdot \left(w_G(x_a, x_b) \cdot \prod_{ij \in F - ab} d_{ij} \pm \sqrt{\varepsilon} |F| \right)$$

where the sum is over all \mathcal{P} -maps $g: [k] \rightarrow V(G)$ such that $g(a) = x_a$ and $g(b) = x_b$.

Proof. The argument is the same as in the counting lemma, except that edges incident to a, b have to be handled differently. First note that for every $c \in [k] - a - b$ and every fixed $x_b \in V_b$, by Observation 7.9 (with $g(x_c) := w_G(x_c, x_b)$), the following holds for all but at most $2\sqrt{\varepsilon}|V_a|$ vertices x_a in V_a :

$$\sum_{x_c} w_G(x_a, x_c) w_G(x_c, x_b) = d_{ac} \left(\sum_{x_c} w_G(x_c, x_b) \right) \pm \sqrt{\varepsilon}|V_c|. \quad (**)$$

For each $c \in [k] - a - b$ and each $x_b \in V_b$, we will ignore those edges going to $x_a \in V_a$ that fail (**).

Similarly for each $c \in [k] - a - b$, by Observation 7.9 (with $g(x_c) := 1$) the following holds for all but at most $2\sqrt{\varepsilon}|V_b|$ vertices $x_b \in V_b$:

$$\sum_{x_c} w_G(x_b, x_c) = d_{bc}|V_c| \pm \sqrt{\varepsilon}|V_c|. \quad (***)$$

We ignore all edges $x_a x_b \in E(V_a, V_b)$ incident to x_b for which (***) fails. Thus for all but $\leq 2 \cdot k \cdot |V_b| \cdot 2\sqrt{\varepsilon}|V_a|$ edges $x_a x_b \in E(V_a, V_b)$, (**) and (***) hold for all $c \in [k] - a - b$.

Fix any such $x_a \in V_a, x_b \in V_b$. We wish to approximate

$$\begin{aligned} & \sum_{(x_i)_{i \in [k] - a - b}} \prod_{ij \in F} w_G(x_i, x_j) = \\ & w_G(x_a, x_b) \cdot \sum_{(x_i)_{i \in [k] - a - b}} \prod_{ij \in F - ab} w_G(x_i, x_j) = \dots \end{aligned}$$

Just as in the proof of the counting lemma, we replace factors $w_G(x_i, x_j)$ by d_{ij} one by one. We first do this for pairs in $F_0 := \{ij \in F \mid i, j \neq a, b\}$, since the argument works without change, incurring an error of $\pm \varepsilon |F_0| \prod_{i \in [k] - a - b} |V_i|$ (we denote this by \simeq for simplicity and sum up the errors at the end of the proof). Since d_{ij} does not depend on the choice of $x_i \in V_i, x_j \in V_j$, we can rearrange:

$$\dots \simeq w_G(x_a, x_b) \left(\prod_{ij \in F_0} d_{ij} \right) \cdot \sum_{(x_i)_{i \in [k] - a - b}} \prod_{i \in [k] - a - b} w_G(x_a, x_i) w_G(x_i, x_b) = \dots$$

Then, for each $c \in [k] - a - b$ we can replace ac by isolating the factors that depend on x_c and applying (**) (as before $w'_G(x_a, x_i)$ denotes either $w_G(x_a, x_i)$ or d_{ai} depending on whether we have already replaced ai):

$$\begin{aligned} \dots &= w_G(x_a, x_b) \left(\prod_{ij \in F_0} d_{ij} \right) \cdot \sum_{x_c} w_G(x_a, x_c) w_G(x_c, x_b) \cdot \\ & \sum_{(x_i)_{i \in [k] - a - b - c}} \prod_{i \in [k] - a - b - c} w'_G(x_a, x_i) w_G(x_i, x_b) \simeq \dots \end{aligned}$$

Having thus replaced all edges ac for $c \in [k] - a - b$, the only remaining edges are of the form

ib for $i \in [k] - a - b$, so by denoting $F_1 := \{ij \in F \mid i, j \neq b\}$ the expression becomes:

$$\begin{aligned} \cdots &\simeq w_G(x_a, x_b) \left(\prod_{ij \in F_1} d_{ij} \right) \cdot \sum_{(x_i)_{i \in [k]-a-b}} \prod_{i \in [k]-a-b} w_G(x_i, x_b) = \\ &w_G(x_a, x_b) \left(\prod_{ij \in F_1} d_{ij} \right) \cdot \prod_{i \in [k]-a-b} \left(\sum_{x_i \in V_i} w_G(x_i, x_b) \right) \simeq \\ &w_G(x_a, x_b) \left(\prod_{ij \in F} d_{ij} \right) \cdot \prod_{i \in [k]-a-b} |V_i|, \end{aligned}$$

where the last approximation follows from (**). For each of the $|F|$ approximations used, the incurred additive error on the whole expression was at most $\pm \sqrt{\varepsilon} \prod_{i \in [k]-a-b} |V_i|$. \square

We are now ready to prove Theorem 7.7. The proof strategy was outlined above: map $G_{/\mathcal{P}}$ to a random copy of F in G , where F marks heavy-enough edges of $G_{/\mathcal{P}}$.

Proof of Theorem 7.7. Let G be a graph with density $c := \frac{E(G)}{n^2}$. Let $\varepsilon_0 < 1$, $\frac{1}{k} \leq \frac{c}{10} \frac{\varepsilon_0}{1+\varepsilon_0}$ and $\varepsilon \leq \left(\frac{1}{k}\right)^{8k^2}$, and suppose G has an ε -regular k -partition $\mathcal{P} = (V_1, \dots, V_k)$. We claim that $d_{\text{opt}}(G, G_{/\mathcal{P}}) \leq \varepsilon_0$. As mentioned above, $G \succeq G_{/\mathcal{P}}$ holds trivially, so it remains to show an overcast from $G_{/\mathcal{P}}$ to $e^{-\varepsilon_0}G$.

Let $d_{ij} := \frac{w_G(V_i, V_j)}{|V_i||V_j|}$ for $ij \in [k]^2$. Let $F \subseteq \binom{[k]}{2}$ be the set of edges ij such that $i \neq j$ and $d_{ij} \geq \frac{1}{k}$. Note that $\prod_{ij \in F} d_{ij} \geq \left(\frac{1}{k}\right)^{|F|} \geq \left(\frac{1}{k}\right)^{k^2} \geq \varepsilon^{1/8}$. Let G' be the subgraph of G obtained by removing:

- $E(G[V_i])$, for $i \in [k]$ (the total weight removed in this step is $\leq k \left(\frac{n}{k}\right)^2$)
- $E_G(V_i, V_j)$, for $ij \notin F$ (their total weight is $\leq k^2 \cdot \frac{1}{k} \cdot \left(\frac{n}{k}\right)^2$)
- edges of weight $< \varepsilon^{1/8}$ (if G is $[0, 1]$ -weighted; their total weight is $\leq \varepsilon^{1/8} n^2$)
- edges $x_a x_b \in V_a \times V_b$ for which the Extension Lemma 7.11 does not hold, for each $ab \in F$ (their total weight is $\leq |F| \cdot 2k\sqrt{\varepsilon} \left(\frac{n}{k}\right)^2$).

The total weight of removed edges is

$$\|G\|_1 - \|G'\|_1 \leq n^2 \left(\frac{1}{k} + \frac{1}{k} + \varepsilon^{1/8} + 2k\sqrt{\varepsilon} \right) \leq n^2 \cdot \frac{5}{k}.$$

By our assumption on k , $\frac{5}{ck} \leq \frac{1}{2} \frac{\varepsilon_0}{1+\varepsilon_0} < 1$. Since $\|G\|_1 \geq cn^2$, $d_1(G, G') \leq \frac{\|G\|_1 - \|G'\|_1}{\min(\|G\|_1, \|G'\|_1)} \leq \frac{n^2 \cdot \frac{5}{k}}{n^2(c - \frac{5}{k})} = \frac{5/ck}{1-5/ck} \leq \varepsilon_0/2$. Therefore, by Lemma 2.11, $G' \succeq e^{-\varepsilon_0/2}G$ (this direction requires only G' to be loopless, which is true because we removed $E(G[V_i])$ for all i). Thus it remains to show that $G_{/\mathcal{P}} \succeq e^{-\varepsilon_0/2}G'$.

We define an overcast ω from $G_{/\mathcal{P}}$ to $(1 - \frac{\varepsilon_0}{1+\varepsilon_0})G'$ as follows: every \mathcal{P} -map $g: [k] \rightarrow V(G)$ is taken with probability proportional to $\text{hom}_g(F, G)$; that is, $\omega(g) := \text{hom}_g(F, G)/N$ where by the Counting Lemma 7.10, the normalisation factor is (using $\prod_{ij \in F} d_{ij} \geq \varepsilon^{1/8}$):

$$N := \sum_g \text{hom}_g(F, G) = \left(\prod_{i \in [k]} |V_i| \right) \left(\prod_{ij \in F} d_{ij} \pm \varepsilon |F| \right) \leq \left(\prod_{i \in [k]} |V_i| \right) \left(\prod_{ij \in F} d_{ij} \right) (1 + \varepsilon^{7/8} k^2).$$

To verify that ω is indeed an overcast, we need to check that for each edge uv of G'

$$\mathbb{E}_{g \sim \omega} w_{G/\mathcal{P}}(g^{-1}(uv)) \geq (1 - \varepsilon') w_G(u, v).$$

Let $a, b \in [k]$ be such that $u \in V_a$ and $v \in V_b$. By the Extension Lemma 7.11, since we removed from G' edges that do not satisfy it and edges with $w_G(u, v) < \varepsilon^{1/8}$, we have:

$$\begin{aligned} N \cdot \mathbb{E}_{g \sim \omega} w_{G/\mathcal{P}}(g^{-1}(uv)) &= \sum_{g: g(a)=u, g(b)=v} \text{hom}_g(F, G) w_{G/\mathcal{P}}(a, b) = \\ &= w_G(V_a, V_b) \sum_{g: g(a)=u, g(b)=v} \text{hom}_g(F, G) = \\ &= d_{ab} |V_a| |V_b| \left(\prod_{i \in [k] - a - b} |V_i| \right) \left(w_G(u, v) \prod_{ij \in F - ab} d_{ij} \pm \sqrt{\varepsilon} |F| \right) \geq \\ &= \left(\prod_{i \in [k]} |V_i| \right) \left(w_G(u, v) \prod_{ij \in F} d_{ij} - \sqrt{\varepsilon} k^2 \right) \geq \\ &= \left(\prod_{i \in [k]} |V_i| \right) w_G(u, v) \left(\prod_{ij \in F} d_{ij} \right) (1 - \varepsilon^{1/4} k^2). \end{aligned}$$

(The last inequality follows from $w_G(u, v) \cdot \prod_{ij \in F} d_{ij} \geq \varepsilon^{1/8} \cdot \varepsilon^{1/8} = \varepsilon^{1/4}$). Dividing by the upper bound on N , we conclude:

$$\mathbb{E}_{g \sim \omega} w_{G/\mathcal{P}}(g^{-1}(uv)) \geq w_G(u, v) \frac{1 - \varepsilon^{1/4} k^2}{1 + \varepsilon^{7/8} k^2}$$

The ratio here can be bounded quite brutally:

$$\geq \frac{1 - \varepsilon^{1/4} k^2}{1 + \varepsilon^{1/4} k^2} \geq 1 - 2\varepsilon^{1/4} k^2 \geq 1 - \frac{2}{k} \geq 1 - \frac{1}{2} \frac{\varepsilon_0}{1 + \varepsilon_0} \geq \frac{1}{1 + \varepsilon_0/2} \geq e^{-\varepsilon_0/2}.$$

This concludes the proof that $G/\mathcal{P} \succeq e^{-\varepsilon_0/2} G'$ and hence $d_{\text{opt}}(G/\mathcal{P}, G) \leq \varepsilon_0$. \square

8 Hardness of approximation

We show that $\text{Max-Hom}(\mathcal{A}, -)$, where \mathcal{A} is the class of all tournaments (orientations of cliques), has no PTAS. This holds under the *Gap Exponential Time Hypothesis* (Gap-ETH) [69, 26] which states that no $2^{o(n)}$ -time algorithm can distinguish between a satisfiable 3SAT formula and one which is not even $(1 - \varepsilon)$ -satisfiable for some constant $\varepsilon > 0$.

In fact we only require the following weaker conjecture:

Conjecture 8.1. There exists an $\varepsilon > 0$ such that given a $\{0, 1\}$ -valued Max-2-CSP instance with k variables and alphabet size n no $f(k) \cdot n^{\mathcal{O}(1)}$ time algorithm can distinguish between the following two cases:

- there is an assignment satisfying every constraint;
- no assignment satisfies more than $(1 - \varepsilon)$ constraints.

Gap-ETH implies Conjecture 8.1: this follows from a proof by Chalermsook et al. [12], in fact with a much larger approximation gap, which was further improved by Dinur and Manurangsi [28]. Direct proofs for the above simpler version can be found in [65] and [9, Appendix A]. Lokshtanov et al. [65] moreover propose the Parameterised Inapproximability Hypothesis, stating that the above promise problem is $W[1]$ -hard.

The problem can be rephrased as a minor variation of Densest- k -Subgraph (sometimes known as Maximum Colored Subgraph Isomorphism):

Observation 8.2. *Conjecture 8.1 is equivalent to the following. There is an $\varepsilon > 0$ such that no $f(k) \cdot n^{O(1)}$ time algorithm can, given k , a graph G on n vertices, and a proper k -colouring c of it, distinguish between the following two cases:*

- G contains a k -clique v_1, \dots, v_k (without loss of generality $c(v_i) = i$);
- every k -tuple v_1, \dots, v_k with $c(v_i) = i$ induces a subgraph on $< (1 - \varepsilon) \binom{k}{2}$ edges in G .

(Indeed, the k variables in the Max-2-CSP correspond to v_1, \dots, v_k , the set of vertices coloured i is the alphabet for variable v_i , and the edges between two colour sets define a constraint). As a side note, we remark that an inspection of the proof of [12, Theorem 5.12] gives that Gap-ETH implies that the above is hard even if the soundness case is strengthened as follows, for *any* constant $\delta > 0$:

- every k -tuple v_1, \dots, v_k (regardless of colours) induces a subgraph on $< \delta \binom{k}{2}$ edges in G .

The problem in Observation 8.2 is almost a maximum graph homomorphism problem on cliques, except that, crucially, the mapping $i \mapsto v_i$ is forced to be injective. To show that $\text{Max-Hom}(\mathcal{A}, -)$ is hard for the class \mathcal{A} of tournaments, intuitively, we use the fact that a map from a random tournament to itself must map most arcs to themselves (and is hence approximately injective). This is formalised as follows.

Lemma 8.3. *For every $\delta > 0$, there exists constants $0 < \lambda < \delta$ and $N \geq 1$ such that the following holds. For every $k \geq N$, there is an orientation \mathbb{A} of the clique of size k such that every mapping $g : A \rightarrow A$ of \mathbb{A} to itself with $\text{value}(g) \geq (1 - \lambda) \binom{k}{2}$ must map at least $(1 - \delta) \binom{k}{2}$ arcs to themselves.*

Proof. For $\delta > 0$, denote $m := \binom{k}{2}$ and choose $N \geq 1$ such that $k \log_2 k \leq \frac{\delta}{3}m$, for all $k \geq N$. Let $\lambda > 0$ be constant to be chosen later. Let \mathbb{A} be a random orientation of the clique of size k with $k \geq N$ (each edge is independently oriented in either direction with probability $\frac{1}{2}$). We will show that with positive probability \mathbb{A} admits no map $g : A \rightarrow A$ to itself with the properties that $\text{value}(g) \geq (1 - \lambda)m$ but less than $(1 - \delta) \binom{k}{2}$ arcs of \mathbb{A} are mapped identically by g .

If a map as above existed, it would imply the existence of a set F of arcs of \mathbb{A} with $|F| \leq \lambda m$ and a mapping $g : A \rightarrow A$ such that g maps all the arcs of $\mathbb{A} - F$ correctly, and such that g maps less than $(1 - \delta) \binom{k}{2}$ vertex pairs identically. Let us bound the probability that there exist such F, g . The number of possible F is $\leq \sum_{i=0}^{\lambda m} \binom{m}{i} \leq 2^{H(\lambda)m}$. The number of possible g is $\leq k^k$. For fixed F, g , if g maps less than $(1 - \delta) \binom{k}{2}$ vertex pairs identically, then the number of remaining arcs of $\mathbb{A} - F$ is at least $(1 - \lambda)m - (1 - \delta) \binom{k}{2} = (\delta - \lambda)m$; the probability that all these arcs are mapped correctly by g is at most $\frac{1}{2}^{(\delta - \lambda)m/2}$ (each of these

arcs is mapped correctly with probability $\frac{1}{2}$; since the function g forms cycles on the set of arcs, the events for individual arcs are not independent, but if we ignore one arc from each cycle they are; since cycles have length at least 2, we ignore at most $\frac{1}{2}$ of these arcs). Hence in total the probability that some such F, g exist is at most

$$2^{H(\lambda)m} \cdot k^k \cdot 2^{-(\delta-\lambda)m/2} = 2^{-(\frac{\delta}{2}-\frac{\lambda}{2}-H(\lambda))m} \cdot 2^{k \log_2 k} \leq 2^{-(\frac{\delta}{2}-\frac{\lambda}{2}-H(\lambda)-\frac{\delta}{3})m}.$$

This is less than 1 by taking λ small enough so that $\frac{\delta}{6} - \frac{\lambda}{2} - H(\lambda) > 0$. \square

This allows us to make the reduction.

Lemma 8.4. *For every $\delta > 0$, there exists constants $0 < \lambda < \delta$ and $N \geq 1$ such that the following holds. Given $k \geq N$, a graph G on n vertices, and a proper k -colouring c of G , we can compute in $f(k) \cdot n^{\mathcal{O}(1)}$ time an orientation \mathbb{A} of the clique of size k and a directed graph \mathbb{B} such that*

- if G contains a clique of size k , then $\text{opt}(\mathbb{A}, \mathbb{B}) = \binom{k}{2}$,
- if every v_1, \dots, v_k in G with $c(v_i) = i$ induce $< (1 - 2\delta)\binom{k}{2}$ edges, then $\text{opt}(\mathbb{A}, \mathbb{B}) < (1 - \lambda)\binom{k}{2}$.

Proof. For $\delta > 0$, let λ and N be as in Lemma 8.3. Given $k \geq N$, G and a proper k -colouring $c : V(G) \rightarrow \{1, \dots, k\}$ of G , we start by computing an orientation \mathbb{A} of the clique on the set of colours $\{1, \dots, k\}$ as in Lemma 8.3 (in time depending on k only). The directed graph \mathbb{B} has vertex set $V(G)$ and (u, v) is an arc in \mathbb{B} iff $\{u, v\} \in E(G)$ and $(c(u), c(v))$ is an arc in \mathbb{A} . Suppose that G contains a clique $\{v_1, \dots, v_k\}$ of size k . Without loss of generality $c(v_i) = i$. Then $\text{opt}(\mathbb{A}, \mathbb{B}) = \binom{k}{2}$ via the mapping $h(i) := v_i$.

Assume now that $\text{opt}(\mathbb{A}, \mathbb{B}) \geq (1 - \lambda)\binom{k}{2}$, so there is a mapping $g : A \rightarrow B$ with $\text{value}(g) \geq (1 - \lambda)\binom{k}{2}$. Note that $c : B \rightarrow A$ is a homomorphism from \mathbb{B} to \mathbb{A} . It follows that the mapping $c \circ g : A \rightarrow A$ from \mathbb{A} to itself has $\text{value}(c \circ g) \geq (1 - \lambda)\binom{k}{2}$. By Lemma 8.3, we have that $c \circ g$ maps at least $(1 - \delta)\binom{k}{2}$ arcs to themselves. Let F be the set of arcs that are not mapped to themselves by $c \circ g$ (so $|F| \leq \delta\binom{k}{2}$). Let F' be the set of arcs of \mathbb{A} that are mapped incorrectly by g (so $|F'| \leq \lambda\binom{k}{2}$). The remaining arcs, $\mathbb{A} - F - F'$, satisfy the following: their number is at least $(1 - \delta - \lambda)\binom{k}{2} \geq (1 - 2\delta)\binom{k}{2}$; they are mapped by g to some arcs in \mathbb{B} and hence to some edge in G ; and if $i \in \{1, \dots, k\}$ is an endpoint of any of these arcs, then $c(g(i)) = i$. We can hence take $v_i := g(i)$ if i is not isolated in $\mathbb{A} - F - F'$ and take an arbitrary v_i with $c(v_i) = i$ otherwise; the resulting k -tuple induces at least $(1 - 2\delta)\binom{k}{2}$ edges in G and satisfies $c(v_i) = i$. \square

From Observation 8.2 (with some $\varepsilon > 0$) and Lemma 8.4 (with $\delta = \frac{\varepsilon}{2}$) we conclude:

Corollary 8.5. *Assuming Conjecture 8.1, there is a constant $\lambda > 0$ such that $\text{Max-Hom}(\mathcal{A}, -)$ for the class of tournaments \mathcal{A} has no $(1 - \lambda)$ -approximation running in time $f(|\mathbb{A}|)(|\mathbb{A}| + |\mathbb{B}|)^{\mathcal{O}(1)}$.*

In particular, assuming Gap-ETH, there is no PTAS (and actually no *FPT approximation scheme*) for $\text{Max-Hom}(\mathcal{A}, -)$.

9 Open questions

Dichotomy Our results, in particular Lemma 1.4, lead us to believe that perhaps the next question has a positive answer.

Question 9.1. Does $\text{Max-}r\text{-CSP}(\mathcal{G})$ admit a PTAS for every r if and only if \mathcal{G} is fractionally-tw-fragile?

A concrete open question concerns cographs, i.e., graphs of clique-width two: Are cographs pliable? More generally, are graphs of bounded clique-width pliable? Note that while certain techniques and results transfer from structurally sparse graphs to dense graphs (clique-width being an example), this fails for pliability as pliability is not closed under complementation since all classes of dense graphs are pliable.

Some example cases where it would be important to show hardness of approximation (or at least integrality gaps for constant levels of the Sherali-Adams hierarchy) in order to shed light on Question 9.1 are classes of unbounded average degree or classes of 3-regular graphs with unbounded girth. In fact, we do not know of any examples of non-pliable classes of structures \mathcal{A} for which $\text{Max-Hom}(\mathcal{A}, -)$ admits a PTAS and thus it is non inconceivable that pliability captures all tractable cases.

Constant-factor approximation Instead of PTASes one can of course ask about the existence of *some* constant-factor approximation. For fixed signatures, Max-Hom always admits a simple constant-factor approximation: essentially map everything randomly to the densest r -tuple, where r is the maximum arity. For the general $\text{Max-}r\text{-CSP}(\mathcal{G})$ problem the situation is more interesting: in general (when \mathcal{G} is the class of all graphs) a constant-factor approximation is impossible; on the other hand for any monotone class of bounded average degree, there is again a simple solution: because such classes have bounded degeneracy, the edge set can be partitioned into a constant number of trees, where the problem can be solved exactly. The results of [28] imply that if the average degree is too high, the problem is again hard. Can a dichotomy be shown?

Weak hyperfiniteness As shown in Theorem 1.6, monotone hyperfinite classes are fractionally-tw-fragile and have bounded degree. The vertex version of hyperfiniteness is called “weakly hyperfinite” in [73] or “fragmentable” in [37, 36]. It is strictly weaker: stars satisfy it, despite having unbounded degree. Nešetřil and Ossona de Mendez [73] proved that for a monotone class of graphs \mathcal{G} of bounded average degree, \mathcal{G} is weakly hyperfinite if and only if for every $d \in \mathbb{N}$, $\{G \in \mathcal{G} : \max \deg(G) \leq d\}$ is hyperfinite. This suggests a possible extension to graphs of unbounded degree: are monotone weakly hyperfinite classes fractionally-tw-fragile? This would imply a conjecture of Dvořák [29], that all graph classes with strongly sublinear separators are fractionally-tw-fragile. However, it is not even known whether all monotone weakly hyperfinite classes have bounded average degree.

Efficient PTAS As mentioned in the introduction, both dense graphs and hyperfinite graphs can be approximated by constant-size descriptions, and in fact by constant-size random samples. Since size-pliability also approximates with constant-size descriptions, this suggests there may be a general way to sample from such structures to give constant-time approximations (for an appropriate input model). In particular, can property-testing results

for hyperfinite graphs be extended to fractionally-tw-fragile graphs? Perhaps this could be a way to obtain EPTASes¹³ for Max- r -CSPs with fixed alphabets. Our methods seems unlikely to give an EPTAS directly. The analogous question in the exact setting is as follows; we believe it to be open.

Question 9.2. Is 3-colouring fixed-parameter tractable when parameterised by the treewidth of the core (the smallest homomorphically equivalent subgraph)? That is, given a graph G which is promised to have a core of treewidth at most k , can we decide its 3-colourability in time $f(k)|G|^{\mathcal{O}(1)}$ for some function f ?

In this question treewidth could also be replaced by size, in which case an algorithm with running time $\mathcal{O}(n)^k$ is trivial (test every k -subgraph for 3-colourability). A similar algorithm for treewidth is due to Dalmau et al. [16], see also [47, Theorem 3.1].

Acknowledgements

The authors would like to thank the referees of the extended abstract [80] and (in particular) of this full version for many detailed and thoughtful comments.

References

- [1] Noga Alon, Wenceslas Fernandez de la Vega, Ravi Kannan, and Marek Karpinski. Random sampling and approximation of MAX-CSPs. *J. Comput. Syst. Sci.*, 67(2):212–243, 2003. [doi:10.1016/S0022-0000\(03\)00008-4](https://doi.org/10.1016/S0022-0000(03)00008-4).
- [2] Noga Alon, Paul D. Seymour, and Robin Thomas. A separator theorem for graphs with an excluded minor and its applications. In *Proc. 22nd Annual ACM Symposium on Theory of Computing (STOC’90)*, pages 293–299. ACM, 1990. [doi:10.1145/100216.100254](https://doi.org/10.1145/100216.100254).
- [3] Gunnar Andersson and Lars Engebretsen. Property testers for dense constraint satisfaction programs on finite domains. *Random Struct. Algorithms*, 21(1):14–32, 2002. [doi:10.1002/rsa.10041](https://doi.org/10.1002/rsa.10041).
- [4] Sanjeev Arora, David R. Karger, and Marek Karpinski. Polynomial Time Approximation Schemes for Dense Instances of NP-Hard Problems. *J. Comput. Syst. Sci.*, 58(1):193–210, 1999. [doi:10.1006/jcss.1998.1605](https://doi.org/10.1006/jcss.1998.1605).
- [5] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998. [doi:10.1145/278298.278306](https://doi.org/10.1145/278298.278306).
- [6] Brenda S. Baker. Approximation Algorithms for NP-Complete Problems on Planar Graphs. *J. ACM*, 41(1):153–180, 1994. [doi:10.1145/174644.174650](https://doi.org/10.1145/174644.174650).
- [7] Itai Benjamini, Oded Schramm, and Asaf Shapira. Every minor-closed property of sparse graphs is testable. In *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC’08)*, pages 393–402. ACM, 2008. [doi:10.1145/1374376.1374433](https://doi.org/10.1145/1374376.1374433).
- [8] Eli Berger, Zdeněk Dvořák, and Sergey Norin. Treewidth of grid subsets. *Comb.*, 38(6):1337–1352, 2018. [arXiv:1512.06441](https://arxiv.org/abs/1512.06441), [doi:10.1007/s00493-017-3538-5](https://doi.org/10.1007/s00493-017-3538-5).
- [9] Arnab Bhattacharyya, Suprovat Ghoshal, Karthik C. S., and Pasin Manurangsi. Parameterized intractability of Even Set and Shortest Vector Problem from Gap-ETH. Technical report, 2018. [arXiv:1803.09717](https://arxiv.org/abs/1803.09717).

¹³Efficient PTAS (EPTAS) is a PTAS whose running time is $\mathcal{O}(n^c)$ for a constant c independent of ε .

- [10] Jacek Brodzki, Graham A Niblo, Jan Spakula, Rufus Willett, and Nick Wright. Uniform local amenability. *Journal of Noncommutative Geometry*, 7(2):583–604, 2013. [arXiv:1203.6169](#), [doi:10.4171/JNCG/128](#).
- [11] Clément Carbonnel, Miguel Romero, and Stanislav Živný. The complexity of general-valued constraint satisfaction problems seen from the other side. *SIAM J. Comput.*, 51(1):19–69, 2022. [arXiv:1710.03148](#), [doi:10.1137/19m1250121](#).
- [12] Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi, Danupon Nanongkai, and Luca Trevisan. From gap-exponential time hypothesis to fixed parameter tractable inapproximability: Clique, dominating set, and more. *SIAM J. Comput.*, 49(4):772–810, 2020. [arXiv:1708.04218](#), [doi:10.1137/18M1166869](#).
- [13] Moses Charikar, MohammadTaghi Hajiaghayi, and Howard J. Karloff. Improved approximation algorithms for label cover problems. *Algorithmica*, 61(1):190–206, 2011. [doi:10.1007/s00453-010-9464-3](#).
- [14] Artur Czumaj, Asaf Shapira, and Christian Sohler. Testing hereditary properties of nonexpanding bounded-degree graphs. *SIAM J. Comput.*, 38(6):2499–2510, 2009. [doi:10.1137/070681831](#).
- [15] David P. Dailey. Uniqueness of colorability and colorability of planar 4-regular graphs are NP-complete. *Discret. Math.*, 30(3):289–293, 1980. [doi:10.1016/0012-365X\(80\)90236-8](#).
- [16] Víctor Dalmau, Phokion G. Kolaitis, and Moshe Y. Vardi. Constraint satisfaction, bounded treewidth, and finite-variable logics. In *Proc. 8th International Conference on Principles and Practice of Constraint Programming (CP’02)*, volume 2470 of *LNCS*, pages 310–326. Springer, 2002. [doi:10.1007/3-540-46135-3_21](#).
- [17] Anuj Dawar, Martin Grohe, Stephan Kreutzer, and Nicole Schweikardt. Approximation Schemes for First-Order Definable Optimisation Problems. In *Proc. 21st IEEE Symposium on Logic in Computer Science (LICS’06)*, pages 411–420. IEEE Computer Society, 2006. [doi:10.1109/LICS.2006.13](#).
- [18] Wenceslas Fernandez de la Vega. MAX-CUT has a randomized approximation scheme in dense graphs. *Random Struct. Algorithms*, 8(3):187–198, 1996.
- [19] Wenceslas Fernandez de la Vega and Marek Karpinski. A polynomial time approximation scheme for subdense MAX-CUT. Technical Report 044, 2002. URL: <http://eccc.hpi-web.de/eccc-reports/2002/TR02-044/index.html>.
- [20] Wenceslas Fernandez de la Vega and Marek Karpinski. Approximation complexity of nondense instances of MAX-CUT. Technical Report 101, 2006. URL: <http://eccc.hpi-web.de/eccc-reports/2006/TR06-101/index.html>.
- [21] Erik D. Demaine, Fedor V. Fomin, Mohammad Taghi Hajiaghayi, and Dimitrios M. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and H -minor-free graphs. *J. ACM*, 52(6):866–893, 2005. [doi:10.1145/1101821.1101823](#).
- [22] Erik D. Demaine and Mohammad Taghi Hajiaghayi. Bidimensionality: new connections between FPT algorithms and ptass. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’05)*, pages 590–601. SIAM, 2005. URL: <http://dl.acm.org/citation.cfm?id=1070432.1070514>.
- [23] Erik D. Demaine, Mohammad Taghi Hajiaghayi, and Ken-ichi Kawarabayashi. Algorithmic Graph Minor Theory: Decomposition, Approximation, and Coloring. In *Proc. 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*, pages 637–646. IEEE Computer Society, 2005. [doi:10.1109/SFCS.2005.14](#).
- [24] Matt DeVos, Guoli Ding, Bogdan Oporowski, Daniel P. Sanders, Bruce Reed, Paul Seymour, and Dirk Vertigan. Excluding any graph as a minor allows a low tree-width 2-coloring. *J. Comb. Theory B*, 91(1):25–41, 2004. [doi:10.1016/j.jctb.2003.09.001](#).

- [25] Reinhard Diestel. *Graph Theory*. Springer, fourth edition, 2010.
- [26] Irit Dinur. Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover. Technical Report 128, 2016. URL: <https://eccc.weizmann.ac.il/report/2016/128/>.
- [27] Irit Dinur, Eldar Fischer, Guy Kindler, Ran Raz, and Shmuel Safra. PCP characterizations of NP: toward a polynomially-small error-probability. *Comput. Complex.*, 20(3):413–504, 2011. doi:10.1007/s00037-011-0014-4.
- [28] Irit Dinur and Pasin Manurangsi. ETH-Hardness of approximating 2-CSPs and Directed Steiner Network. In *Proc. 9th Innovations in Theoretical Computer Science Conference (ITCS'18)*, volume 94 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 36:1–36:20. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018. arXiv:1805.03867, doi:10.4230/LIPIcs.ITCS.2018.36.
- [29] Zdeněk Dvořák. Sublinear separators, fragility and subexponential expansion. *Eur. J. Comb.*, 52(Part A):103–119, 2016. arXiv:1404.7219, doi:10.1016/j.ejc.2015.09.001.
- [30] Zdeněk Dvořák. Baker game and polynomial-time approximation schemes. In *Proc. 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'20)*, pages 2227–2240. SIAM, 2020. arXiv:1901.01797, doi:10.1137/1.9781611975994.137.
- [31] Zdeněk Dvořák, Tony Huynh, Gwenael Joret, Chun-Hung Liu, and David R. Wood. *Notes on Graph Product Structure Theory*, pages 513–533. Springer International Publishing, 2021. arXiv:2001.08860, doi:10.1007/978-3-030-62497-2_32.
- [32] Zdeněk Dvořák and Jean-Sébastien Sereni. On fractional fragility rates of graph classes. *Electron. J. Comb.*, 27(4):4, 2020. arXiv:1907.12634, doi:10.37236/8909.
- [33] Zdeněk Dvořák and Abhiruk Lahiri. Approximation Schemes for Bounded Distance Problems on Fractionally Treewidth-Fragile Graphs. In *Proc. 29th Annual European Symposium on Algorithms (ESA'21)*, volume 204 of *LIPIcs*, pages 40:1–40:10. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. arXiv:2105.01780, doi:10.4230/LIPIcs.ESA.2021.40.
- [34] Zdeněk Dvořák. Thin graph classes and polynomial-time approximation schemes. In *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'18)*, pages 1685–1701, 2018. arXiv:1704.00125, doi:10.1137/1.9781611975031.110.
- [35] Zdeněk Dvořák and Sergey Norin. Strongly sublinear separators and polynomial expansion. *SIAM J. Discrete Math.*, 30(2):1095–1101, 2016. doi:10.1137/15M1017569.
- [36] Keith Edwards and Graham Farr. *Graph fragmentability*, pages 203–218. Encyclopedia of Mathematics and Its Applications. Cambridge University Press, 2012.
- [37] Keith Edwards and Colin McDiarmid. New upper bounds on harmonious colorings. *J. Graph Theory*, 18(3):257–267, 1994. doi:10.1002/jgt.3190180305.
- [38] Gábor Elek. L2-spectral invariants and convergent sequences of finite graphs. *J. Funct. Anal.*, 10(254):2667–2689, 2008. arXiv:0709.1261, doi:10.1016/j.jfa.2008.01.010.
- [39] Gábor Elek. Uniform local amenability implies Property A. *Proceedings of the American Mathematical Society*, 149(6):2573–2577, 2021. arXiv:1912.00806, doi:10.1090/proc/15387.
- [40] David Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27(3-4):275–291, 2000. doi:10.1007/s004530010020.
- [41] Alan M. Frieze and Ravi Kannan. The regularity lemma and approximation schemes for dense problems. In *Proc. 37th Annual Symposium on Foundations of Computer Science (FOCS'96)*, pages 12–20. IEEE Computer Society, 1996. doi:10.1109/SFCS.1996.548459.

- [42] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995. doi:10.1145/227683.227684.
- [43] Oded Goldreich. *Introduction to property testing*. Cambridge University Press, 2017. URL: <http://www.wisdom.weizmann.ac.il/~oded/pt-intro.html>.
- [44] Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *J. ACM*, 45(4):653–750, 1998. doi:10.1145/285055.285060.
- [45] Alexander Grigoriev and Hans L. Bodlaender. Algorithms for graphs embeddable with few crossings per edge. *Algorithmica*, 49(1):1–11, 9 2007. doi:10.1007/s00453-007-0010-x.
- [46] Martin Grohe. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica*, 23(4):613–632, 2003. doi:10.1007/s00493-003-0037-9.
- [47] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM*, 54(1):1–24, 2007. doi:10.1145/1206035.1206036.
- [48] Martin Grohe, Thomas Schwentick, and Luc Segoufin. When is the evaluation of conjunctive queries tractable? In *Proc. 33rd Annual ACM Symposium on Theory of Computing (STOC’01)*, pages 657–666. ACM, 2001. doi:10.1145/380752.380867.
- [49] Venkatesan Guruswami and Sanjeev Khanna. On the hardness of 4-coloring a 3-colorable graph. *SIAM J. Discrete Math.*, 18(1):30–40, 2004. doi:10.1137/S0895480100376794.
- [50] F. Hadlock. Finding a maximum cut of a planar graph in polynomial time. *SIAM J. Comput.*, 4(3):221–225, 1975. doi:10.1137/0204019.
- [51] Avinatan Hassidim, Jonathan A. Kelner, Huy N. Nguyen, and Krzysztof Onak. Local graph partitions for approximation and testing. In *Proc. 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS’09)*, pages 22–31, 2009. doi:10.1109/FOCS.2009.77.
- [52] Pavol Hell and Jaroslav Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.
- [53] Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and vlsi. *J. ACM*, 32(1):130–136, 1 1985. doi:10.1145/2455.214106.
- [54] Harry B. Hunt III, Riko Jacob, Madhav V. Marathe, and Richard E. Stearns. Towards syntactic characterizations of approximation schemes via predicate and graph decompositions. In *Proc. IEEE Symposium on Foundations of Computer Science (FOCS’98)*, 12 1998. URL: <https://www.osti.gov/biblio/677169>.
- [55] Harry B. Hunt III, Madhav V. Marathe, Venkatesh Radhakrishnan, S.S Ravi, Daniel J. Rosenkrantz, and Richard E. Stearns. NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs. *J. Algorithms*, 26(2):238 – 274, 1998. doi:10.1006/jagm.1997.0903.
- [56] Richard M. Karp. Reducibility Among Combinatorial Problems. In *Complexity of Computer Computations*, pages 85–103, 1972. URL: <http://www.cs.berkeley.edu/%7Eluca/cs172/karp.pdf>.
- [57] Sanjeev Khanna and Rajeev Motwani. Towards a syntactic characterization of PTAS. In *Proc. 28th Annual ACM Symposium on the Theory of Computing (STOC’96)*, pages 329–337, 1996. doi:10.1145/237814.237979.
- [58] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proc. 34th Annual ACM Symposium on Theory of Computing (STOC’02)*, pages 767–775. ACM, 2002. doi:10.1145/509907.510017.
- [59] Alexandr V. Kostochka. Lower bound of the hadwiger number of graphs by their average degree. *Comb.*, 4(4):307–316, 1984. doi:10.1007/BF02579141.

- [60] Robert Krauthgamer and James R. Lee. The intrinsic dimensionality of graphs. *Comb.*, 27(5):551–585, 2007. doi:10.1007/s00493-007-2183-y.
- [61] Akash Kumar, C. Seshadhri, and Andrew Stolman. Random walks and forbidden minors II: a $\text{poly}(de^{-1})$ -query tester for minor-closed properties of bounded degree graphs. In *Proc. 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC'19)*, pages 559–567. ACM, 2019. doi:10.1145/3313276.3316330.
- [62] Michael Langberg, Yuval Rabani, and Chaitanya Swamy. Approximation algorithms for graph homomorphism problems. In *Proc. Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM'06)*, volume 4110 of *Lecture Notes in Computer Science*, pages 176–187. Springer, 2006. doi:10.1007/11830924_18.
- [63] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36(2):177–189, 1979. doi:10.1137/0136016.
- [64] Richard J. Lipton and Robert Endre Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980. doi:10.1137/0209046.
- [65] Daniel Lokshtanov, M. S. Ramanujan, Saket Saurabh, and Meirav Zehavi. Parameterized complexity and approximability of directed odd cycle transversal. In *Proc. 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'20)*, pages 2181–2200. SIAM, 2020. arXiv:1704.04249, doi:10.1137/1.9781611975994.134.
- [66] László Lovász. *Large Networks and Graph Limits*, volume 60 of *Colloquium Publications*. American Mathematical Society, 2012. URL: <http://web.cs.elte.hu/~lovasz/bookxx/hombook-almost.final.pdf>.
- [67] László Lovász and Balázs Szegedy. Szemerédi’s lemma for the analyst. *Geom. Funct. Anal. (GAFA)*, 17(1):252–270, 2007. doi:10.1007/s00039-007-0599-6.
- [68] Pasin Manurangsi and Dana Moshkovitz. Approximating dense max 2-CSPs. In *Proc. Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM'15)*, pages 396–415. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015. arXiv:1507.08348, doi:10.4230/LIPIcs.APPROX-RANDOM.2015.396.
- [69] Pasin Manurangsi and Prasad Raghavendra. A Birthday Repetition Theorem and Complexity of Approximating Dense CSPs. In *Proc. 44th International Colloquium on Automata, Languages, and Programming (ICALP'17)*, volume 80 of *LIPIcs*, pages 78:1–78:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. arXiv:1607.02986, doi:10.4230/LIPIcs.ICALP.2017.78.
- [70] Madhav V. Marathe, Harry B. Hunt III, and Richard E. Stearns. Level-treewidth property, exact algorithms and approximation schemes. In *Proc. ACM Symposium on Theory of Computing (STOC'97)*, 6 1997. URL: <https://www.osti.gov/biblio/471394>.
- [71] Balázs F. Mezei, Marcin Wrochna, and Stanislav Živný. PTAS for sparse general-valued csps. *ACM Trans. Algorithms*, 19(2):14:1–14:31, 2023. arXiv:2012.12607, doi:10.1145/3569956.
- [72] Guy Moshkovitz and Asaf Shapira. Decomposing a graph into expanding subgraphs. *Random Struct. Algorithms*, 52(1):158–178, 2018. doi:10.1002/rsa.20727.
- [73] Jaroslav Nešetřil and Patrice Ossona de Mendez. *Property Testing, Hyperfiniteness and Separators*, pages 363–379. Springer, 2012. doi:10.1007/978-3-642-27875-4_16.
- [74] Jaroslav Nešetřil and Patrice Ossona de Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *Eur. J. Comb.*, 27(6):1022–1041, 2006. doi:10.1016/j.ejc.2005.01.010.
- [75] Jaroslav Nešetřil and Patrice Ossona de Mendez. Grad and classes with bounded expansion II. algorithmic aspects. *Eur. J. Comb.*, 29(3):777–791, 2008. doi:10.1016/j.ejc.2006.07.014.
- [76] Ilan Newman and Christian Sohler. Every property of hyperfinite graphs is testable. *SIAM J. Comput.*, 42(3):1095–1112, 2013. doi:10.1137/120890946.

- [77] Christos H. Papadimitriou and Mihalis Yannakakis. Optimization, Approximation, and Complexity Classes. *J. Comput. Syst. Sci.*, 43(3):425–440, 1991. doi:10.1016/0022-0000(91)90023-X.
- [78] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proc. 40th Annual ACM Symposium on Theory of Computing (STOC'08)*, pages 245–254, 2008. doi:10.1145/1374376.1374414.
- [79] Vojtěch Rödl and Mathias Schacht. Regular partitions of hypergraphs: Regularity lemmas. *Comb. Probab. Comput.*, 16(6):833–885, 2007. doi:10.1017/S0963548307008553.
- [80] Miguel Romero, Marcin Wrochna, and Stanislav Živný. Treewidth-Pliability and PTAS for Max-CSPs. In *Proc. 2021 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'21)*, pages 473–483. SIAM, 2021. doi:10.1137/1.9781611976465.29.
- [81] Sartaj Sahni and Teofilo F. Gonzalez. P-Complete Approximation Problems. *J. ACM*, 23(3):555–565, 1976. doi:10.1145/321958.321975.
- [82] Alexander Schrijver. *Theory of linear and integer programming*. John Wiley & Sons, Inc., 1986.
- [83] Hanif D. Sherali and Warren P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM J. Discrete Math.*, 3(3):411–430, 1990. doi:10.1137/0403036.
- [84] Joel Spencer. Counting extensions. *J. Comb. Theory, Ser. A*, 55(2):247–255, 1990. doi:10.1016/0097-3165(90)90070-D.
- [85] Endre Szemerédi. Regular partitions of graphs. In *Proc. Colloq. Inter. CNRS*, pages 399–401, 1978. URL: <http://i.stanford.edu/pub/ctr/reports/cs/tr/75/489/CS-TR-75-489.pdf>.
- [86] Johan Thapper and Stanislav Živný. The complexity of finite-valued CSPs. *J. ACM*, 63(4), 2016. Article No. 37. doi:10.1145/2974019.
- [87] Carsten Thomassen. Girth in graphs. *J. Comb. Theory, Ser. B*, 35(2):129–141, 1983. doi:10.1016/0095-8956(83)90067-9.
- [88] Jan van den Heuvel, Patrice Ossona de Mendez, Daniel Quiroz, Roman Rabinovich, and Sebastian Siebertz. On the generalised colouring numbers of graphs that exclude a fixed minor. *Eur. J. Comb.*, 66:129–144, 2017. arXiv:1602.09052, doi:10.1016/j.ejc.2017.06.019.

A Farkas’ lemma and proofs of Proposition 2.5 and Lemma 4.3

Farkas’ lemma is the fundamental duality for systems of linear inequalities.

Lemma A.1 (Farkas’ lemma [82, Corollary 7.1d]). *Let A be an $m \times n$ rational matrix and $\bar{b} \in \mathbb{Q}^m$. Then, exactly one of the two holds:*

- $A\bar{x} = \bar{b}$ for some $\bar{x} \in \mathbb{Q}^n$ with $\bar{x} \geq 0$, or
- $A^T \bar{y} \geq 0$ and $\bar{b}^T \bar{y} < 0$ for some $\bar{y} \in \mathbb{Q}^m$.

For the duality between the existence of overcasts and the overcast relation \succeq , we use Farkas’ lemma in the following form:

Lemma A.2 (Farkas’ lemma, variant 1). *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 $i = 1, \dots, n$.

Proof. The first condition is equivalent to the existence of a solution in variables $x_i \in \mathbb{Q}_{\geq 0}$ ($i = 1, \dots, n$) and $s_j \in \mathbb{Q}_{\geq 0}$ ($j = 1, \dots, m$) of the following system:

$$\begin{aligned} \sum_i x_i &= 1 \\ \sum_i A_{i,j} x_i - s_j &= b_j \quad (\text{for } j = 1, \dots, m). \end{aligned}$$

By Lemma A.1, this system has a solution if and only if the following system has no solution in variables $z \in \mathbb{Q}$ and $y_j \in \mathbb{Q}^m$ for $j = 1, \dots, m$:

$$\begin{aligned} z + \sum_j A_{i,j} y_j &\geq 0 \quad (\text{for } i = 1, \dots, n) \\ -y_j &\geq 0 \quad (\text{for } j = 1, \dots, m) \\ z + \sum_j b_j y_j &< 0 \end{aligned}$$

Equivalently, there are no $y'_j = -y_j \in \mathbb{Q}_{\geq 0}^m$ such that $\sum_j A_{i,j} y'_j \leq z < \sum_j b_j y'_j$ (for $i = 1, \dots, n$). \square

Proposition (Proposition 2.5 restated). *Let \mathbb{A} and \mathbb{B} be σ -structures. Then, $\mathbb{A} \succeq \mathbb{B}$ if and only if there is an overcast from \mathbb{A} to \mathbb{B} .*

Proof. First, suppose that there exists an overcast ω from \mathbb{A} to \mathbb{B} . Let \mathbb{C} be a σ -structure. Then, if h is a maximum-value mapping from \mathbb{B} to \mathbb{C} we have

$$\begin{aligned} \text{opt}(\mathbb{B}, \mathbb{C}) &= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B})} f^{\mathbb{B}}(\mathbf{x}) f^{\mathbb{C}}(h(\mathbf{x})) \leq \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B})} \left(\sum_{g \in B^A} \omega(g) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) \right) f^{\mathbb{C}}(h(\mathbf{x})) \\ &= \sum_{g \in B^A} \omega(g) \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B})} f^{\mathbb{A}}(g^{-1}(\mathbf{x})) f^{\mathbb{C}}(h(\mathbf{x})) \right) \\ &= \sum_{g \in B^A} \omega(g) \left(\sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{y}) f^{\mathbb{C}}(h(g(\mathbf{y}))) \right) \end{aligned}$$

and hence there exists $g \in B^A$ such that $\text{opt}(\mathbb{B}, \mathbb{C}) \leq \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A})} f^{\mathbb{A}}(\mathbf{y}) f^{\mathbb{C}}(h(g(\mathbf{y}))) = \text{value}(h \circ g) \leq \text{opt}(\mathbb{A}, \mathbb{C})$. Therefore, $\mathbb{A} \succeq \mathbb{B}$. For the converse implication, we shall use Lemma A.2. If there is no overcast from \mathbb{A} to \mathbb{B} , this means there are no numbers $\omega(g) \in \mathbb{Q}_{\geq 0}$ (for $g \in B^A$) such that $\sum_g \omega(g) = 1$ and $\sum_{g \in B^A} \omega(g) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) \geq f^{\mathbb{B}}(\mathbf{x})$ for $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})$. By the lemma above, this is equivalent to the existence of $y(f, \mathbf{x}) \in \mathbb{Q}_{\geq 0}$ (for $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})$) such that

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

Let $\mathbb{B}_{\bar{y}}$ be the σ -structure with domain B such that $f^{\mathbb{B}_{\bar{y}}}(\mathbf{x}) = y(f, \mathbf{x})$, for all $(f, \mathbf{x}) \in \text{tup}(\mathbb{B})$ ($f^{\mathbb{B}_{\bar{y}}}(\mathbf{x}) = 0$ otherwise). By the above, $\text{opt}(\mathbb{A}, \mathbb{B}_{\bar{y}}) < \text{opt}(\mathbb{B}, \mathbb{B}_{\bar{y}})$ and hence $\mathbb{A} \not\succeq \mathbb{B}$. \square

For the duality between ε -thin distributions of modulators and weights avoiding any ε -small modulator, we use a variant of Farkas' lemma obtained by taking $-A$ and $-b$ in Lemma A.2.

Lemma A.3 (Farkas' lemma, variant 2). *Let A be an $m \times n$ rational matrix and let $\bar{b} \in \mathbb{Q}^m$. 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 $\sum_i A_{i,j} x_i \leq b_j$ 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 $i = 1, \dots, n$.*

Lemma (Lemma 4.3 restated). *Let \mathcal{F} be a family of subsets of a set V . The following are equivalent:*

- *there is an ε -thin distribution π of sets $X \in \mathcal{F}$ (i.e., for all $v \in V$, $\Pr_{X \sim \pi}[v \in X] \leq \varepsilon$);*
- *for all non-negative weights $(w(v))_{v \in V}$, there is an $X \in \mathcal{F}$ such that $w(X) \leq \varepsilon \cdot w(V)$.*

Proof. The first item is equivalent to the existence of numbers $\pi(X) \in \mathbb{Q}_{\geq 0}$ for $X \in \mathcal{F}$ such that $\sum_X \pi(X) = 1$ and for all $v \in V$, $\sum_X [v \in X] \cdot \pi(X) \leq \varepsilon$. By Lemma A.3, this holds if and only if there are no numbers $w(v) \in \mathbb{Q}_{\geq 0}$ for $v \in V$ such that for all $X \in \mathcal{F}$, $\sum_v [v \in X] \cdot w(v) > \sum_v \varepsilon \cdot w(v)$. \square

B Proof of Proposition 3.3

Let $\text{SA}_k(\mathbb{A}, \mathbb{C})$ denote the Sherali-Adams linear programming relaxation of $\text{Max-Hom}(\mathbb{A}, \mathbb{C})$, given in Figure 1 in Section 3. Recall that $\text{opt}_k(\mathbb{A}, \mathbb{C})$ denotes its optimum value and we write $\mathbb{A} \succeq_k \mathbb{B}$ if $\text{opt}_k(\mathbb{A}, \mathbb{C}) \geq \text{opt}_k(\mathbb{B}, \mathbb{C})$ for all structures \mathbb{C} with the same signature as \mathbb{A} and \mathbb{B} .

Proposition (Proposition 3.3 restated). *Let \mathbb{A} and \mathbb{B} be σ -structures and $k \geq \max_{f \in \sigma} \text{ar } f$. If there is an overcast from \mathbb{A} to \mathbb{B} then $\mathbb{A} \succeq_k \mathbb{B}$.*

Proof. Let \mathbb{C} be an arbitrary σ -structure, ω be an overcast from \mathbb{A} to \mathbb{B} and λ be an optimal solution to $\text{SA}_k(\mathbb{B}, \mathbb{C})$. (Recall that for a tuple \mathbf{x} we denote by $\text{Set}(\mathbf{x})$ the set of elements appearing in \mathbf{x} .) We have that

$$\begin{aligned}
\text{opt}_k(\mathbb{B}, \mathbb{C}) &= \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B}), s: \text{Set}(\mathbf{x}) \rightarrow C} \lambda(\text{Set}(\mathbf{x}), s) f^{\mathbb{B}}(\mathbf{x}) f^{\mathbb{C}}(s(\mathbf{x})) \\
&\leq \sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B}), s: \text{Set}(\mathbf{x}) \rightarrow C} \left(\sum_{g \in B^A} \omega(g) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) \right) \lambda(\text{Set}(\mathbf{x}), s) f^{\mathbb{C}}(s(\mathbf{x})) \\
&= \sum_{g \in B^A} \omega(g) \left(\sum_{(f, \mathbf{x}) \in \text{tup}(\mathbb{B}), s: \text{Set}(\mathbf{x}) \rightarrow C} \lambda(\text{Set}(\mathbf{x}), s) f^{\mathbb{A}}(g^{-1}(\mathbf{x})) f^{\mathbb{C}}(s(\mathbf{x})) \right) \\
&= \sum_{g \in B^A} \omega(g) \left(\sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A}), s: g(\text{Set}(\mathbf{y})) \rightarrow C} \lambda(g(\text{Set}(\mathbf{y})), s) f^{\mathbb{A}}(\mathbf{y}) f^{\mathbb{C}}(s(g(\mathbf{y}))) \right)
\end{aligned}$$

and hence there is $g : A \rightarrow B$ such that

$$\text{opt}_k(\mathbb{B}, \mathbb{C}) \leq \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A}), s: g(\text{Set}(\mathbf{y})) \rightarrow C} \lambda(g(\text{Set}(\mathbf{y})), s) f^{\mathbb{A}}(\mathbf{y}) f^{\mathbb{C}}(s(g(\mathbf{y}))) \quad (1)$$

For $Y \in \binom{A}{\leq k}$ and $r : Y \rightarrow C$, we define

$$\lambda'(Y, r) = \begin{cases} \lambda(g(Y), s) & \text{if there exists } s : g(Y) \rightarrow C \text{ such that } s \circ g = r \\ 0 & \text{otherwise} \end{cases}$$

Note that λ' is a feasible solution of $\text{SA}_k(\mathbb{A}, \mathbb{C})$. Indeed, for $Y \in \binom{A}{\leq k}$, we have

$$\sum_{r: Y \rightarrow C} \lambda'(Y, r) = \sum_{s: g(Y) \rightarrow C} \lambda'(Y, s \circ g) = \sum_{s: g(Y) \rightarrow C} \lambda(g(Y), s) = 1.$$

Moreover, let $Z \subseteq Y \in \binom{A}{\leq k}$, and $r : Z \rightarrow C$. If there is no $s : g(Z) \rightarrow C$ such that $s \circ g = r$, then

$$\lambda'(Z, r) = 0 = \sum_{t: Y \rightarrow C, t|_Z = r} \lambda'(Y, t).$$

If such a mapping s exists, then

$$\begin{aligned} \sum_{t: Y \rightarrow C, t|_Z = r} \lambda'(Y, t) &= \sum_{s': g(Y) \rightarrow C, s'|_{g(Z)} = s} \lambda'(Y, s' \circ g) \\ &= \sum_{s': g(Y) \rightarrow C, s'|_{g(Z)} = s} \lambda(g(Y), s') \\ &= \lambda(g(Z), s) \\ &= \lambda'(Z, r). \end{aligned}$$

Since λ' is feasible and by (1), we conclude that

$$\begin{aligned} \text{opt}_k(\mathbb{A}, \mathbb{C}) &\geq \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A}), r: \text{Set}(\mathbf{y}) \rightarrow C} \lambda'(\text{Set}(\mathbf{y}), r) f^{\mathbb{A}}(\mathbf{y}) f^{\mathbb{C}}(r(\mathbf{y})) \\ &= \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A}), s: g(\text{Set}(\mathbf{y})) \rightarrow C} \lambda'(\text{Set}(\mathbf{y}), s \circ g) f^{\mathbb{A}}(\mathbf{y}) f^{\mathbb{C}}(s(g(\mathbf{y}))) \\ &= \sum_{(f, \mathbf{y}) \in \text{tup}(\mathbb{A}), s: g(\text{Set}(\mathbf{y})) \rightarrow C} \lambda(g(\text{Set}(\mathbf{y})), s) f^{\mathbb{A}}(\mathbf{y}) f^{\mathbb{C}}(s(g(\mathbf{y}))) \\ &\geq \text{opt}_k(\mathbb{B}, \mathbb{C}). \end{aligned}$$

□